

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 447 576 A1

(12)

EUROPEAN PATENT APPLICATION(21) Application number: **90105101.1**(51) Int. Cl.⁵: **G06F 11/18**(22) Date of filing: **19.03.90**(43) Date of publication of application:
25.09.91 Bulletin 91/39(64) Designated Contracting States:
BE CH DK ES LI LU NL(71) Applicant: **TANDEM COMPUTERS
INCORPORATED**
10435 N. Tantau Avenue
Cupertino, California 95014-0709(US)(72) Inventor: **Cutts, Richard W, Jr.**
1312 Elm Street
Georgetown, Texas 78626(US)
Inventor: **Jewett, Douglas E.**
12401 Wycliff Lane
Austin, Texas 78727(US)
Inventor: **Southworth, Richard A.**
10008 Brandywine Circle
Austin, Texas 78750(US)


Inventor: **Debacker Kenneth C.**
15702 Monona Cove
Austin, Texas 78717(US)
Inventor: **Mehta, Nikhil A.**
1715 Prairie Hen Cove
Austin, Texas 78758(US)
Inventor: **Allison, John David**
1000 Hawk Court
Round Rock, Texas 78681(US)
Inventor: **Horst, Robert W.**
2804 Robeson Park Dr.
Champaign, Illinois 61821(US)

(74) Representative: **Altenburg, Udo, Dipl.-Phys. et
al**
Patent- und Rechtsanwälte
Bardehle-Pagenberg-Dost-Altenburg &
Partner Postfach 86 06 20
W-8000 München 86(DE)(54) **Synchronization of fault-tolerant computer system having multiple processors.**

(57) A computer system in a fault-tolerant configuration employs three identical CPUs executing the same instruction stream, with two identical, self-checking memory modules storing duplicates of the same data. Memory references by the three CPUs are made by three separate busses connected to three separate ports of each of the two memory modules. The three CPUs are loosely synchronized, as by detecting events such as memory references and stalling any CPU ahead of others until all ex-

ecute the function simultaneously; interrupts can be synchronized by ensuring that all three CPUs implement the interrupt at the same point in their instruction stream. Memory references via the separate CPU-to-memory busses are voted at the three separate ports of each of the memory modules. I/O functions are implemented using two identical I/O busses, each of which is separately coupled to only one of the memory modules. A number of I/O processors are coupled to both I/O busses.

EP 0 447 576 A1



RELATED CASES: This application discloses subject matter also disclosed in copending U.S. patent applications Ser. Nos. 282,469, 282,540, and 282,629, filed December 9, 1988, and Ser. Nos. 283,573 and 283, 574, filed December 13, 1988; further, this application is a continuation-in-part of U.S. patent application Serial No. 118,503, filed November 9, 1987, by Robert W. Horst; all of said applications being assigned to Tandem Computers Incorporated, the assignee of this invention.

BACKGROUND OF THE INVENTION

This invention relates to computer systems, and more particularly synchronizing methods for a fault-tolerant system using multiple CPUs.

Highly reliable digital processing is achieved in various computer architectures employing redundancy. For example, TMR (triple modular redundancy) systems may employ three CPUs executing the same instruction stream, along with three separate main memory units and separate I/O devices which duplicate functions, so if one of each type of element fails, the system continues to operate. Another fault-tolerant type of system is shown in U.S. Patent 4,228,496, issued to Katzman et al, for "Multiprocessor System", assigned to Tandem Computers Incorporated. Various methods have been used for synchronizing the units in redundant systems; for example, in said prior application Ser. No. 118,503, filed Nov. 9, 1987, by R. W. Horst, for "Method and Apparatus for Synchronizing a Plurality of Processors", also assigned to Tandem Computers Incorporated, a method of "loose" synchronizing is disclosed, in contrast to other systems which have employed a lock-step synchronization using a single clock, as shown in U.S. Patent 4,453,215 for "Central Processing Apparatus for Fault-Tolerant Computing", assigned to Stratus Computer, Inc. A technique called "synchronization voting" is disclosed by Davies & Wakerly in "Synchronization and Matching in Redundant Systems", IEEE Transactions on Computers June 1978, pp. 531-539. A method for interrupt synchronization in redundant fault-tolerant systems is disclosed by Yondea et al in Proceeding of 15th Annual Symposium on Fault-Tolerant Computing, June 1985, pp. 246-251, "Implementation of Interrupt Handler for Loosely Synchronized TMR Systems". U.S. Patent 4,644,498 for "Fault-Tolerant Real Time Clock" discloses a triple modular redundant clock configuration for use in a TMR computer system. U.S. Patent 4,733,353 for "Frame Synchronization of Multiply Redundant Computers" discloses a synchronization method using separately-clocked CPUs which are periodically synchronized by executing a synch frame.

As high-performance microprocessor devices

have become available, using higher clock speeds and providing greater capabilities, such as the Intel 80386 and Motorola 68030 chips operating at 25-MHz clock rates, and as other elements of computer systems such as memory, disk drives, and the like have correspondingly become less expensive and of greater capability, the performance and cost of high-reliability processors has been required to follow the same trends. In addition, standardization on a few operating systems in the computer industry in general has vastly increased the availability of applications software, so a similar demand is made on the field of high-reliability systems; i.e., a standard operating system must be available.

It is therefore the principal object of this invention to provide an improved high-reliability computer system, particularly of the fault-tolerant type. Another object is to provide an improved redundant, fault-tolerant type of computing system, and one in which high performance and reduced cost are both possible; particularly, it is preferable that the improved system avoid the performance burdens usually associated with highly redundant systems. A further object is to provide a high-reliability computer system in which the performance, measured in reliability as well as speed and software compatibility, is improved but yet at a cost comparable to other alternatives of lower performance. An additional object is to provide a high-reliability computer system which is capable of executing an operating system which uses virtual memory management with demand paging, and having protected (supervisory or "kernel") mode; particularly an operating system also permitting execution of multiple processes; all at a high level of performance.

SUMMARY OF THE INVENTION

In accordance with one embodiment of the invention, a computer system employs three identical CPUs typically executing the same instruction stream, and has two identical, self-checking memory modules storing duplicates of the same data. A configuration of three CPUs and two memories is therefore employed, rather than three CPUs and three memories as in the classic TMR systems. Memory references by the three CPUs are made by three separate busses connected to three separate ports of each of the two memory modules. In order to avoid imposing the performance burden of fault-tolerant operation on the CPUs themselves, and imposing the expense, complexity and timing problems of fault-tolerant clocking, the three CPUs each have their own separate and independent clocks, but are loosely synchronized, as by detecting events such as memory references and stalling

any CPU ahead of others until all execute the function simultaneously; the interrupts are also synchronized to the CPUs ensuring that the CPUs execute the interrupt at the same point in their instruction stream. The three asynchronous memory references via the separate CPU-to-memory busses are voted at the three separate ports of each of the memory modules at the time of the memory request, but read data is not voted when returned to the CPUs.

The two memories both perform all write requests received from either the CPUs or the I/O busses, so that both are kept up-to-date, but only one memory module presents read data back to the CPUs or I/Os in response to read requests; the one memory module producing read data is designated the "primary" and the other is the back-up. Accordingly, incoming data is from only one source and is not voted. The memory requests to the two memory modules are implemented while the voting is still going on, so the read data is available to the CPUs a short delay after the last one of the CPUs makes the request. Even write cycles can be substantially overlapped because DRAMs used for these memory modules use a large part of the write access to merely read and refresh, then if not strobed for the last part of the write cycle the read is non-destructive; therefore, a write cycle begins as soon as the first CPU makes a request, but does not complete until the last request has been received and voted good. These features of non-voted read-data returns and overlapped accesses allow fault-tolerant operation at high performance, but yet at minimum complexity and expense.

I/O functions are implemented using two identical I/O busses, each of which is separately coupled to only one of the memory modules. A number of I/O processors are coupled to both I/O busses, and I/O devices are coupled to pairs of the I/O processors but accessed by only one of the I/O processors. Since one memory module is designated primary, only the I/O bus for this module will be controlling the I/O processors, and I/O traffic between memory module and I/O is not voted. The CPUs can access the I/O processors through the memory modules (each access being voted just as the memory accesses are voted), but the I/O processors can only access the memory modules, not the CPUs; the I/O processors can only send interrupts to the CPUs, and these interrupts are collected in the memory modules before presenting to the CPUs. Thus synchronization overhead for I/O device access is not burdening the CPUs, yet fault tolerance is provided. If an I/O processor fails, the other one of the pair can take over control of the I/O devices for this I/O processor by merely changing the addresses used for the I/O device in the I/O page table maintained by the operating system. In

this manner, fault tolerance and reintegration of an I/O device is possible without system shutdown, and yet without hardware expense and performance penalty associated with voting and the like in these I/O paths.

The memory system used in the illustrated embodiment is hierarchical at several levels. Each CPU has its own cache, operating at essentially the clock speed of the CPU. Then each CPU has a local memory not accessible by the other CPUs, and virtual memory management allows the kernel of the operating system and pages for the current task to be in local memory for all three CPUs, accessible at high speed without fault-tolerance overhead such as voting or synchronizing imposed. Next is the memory module level, referred to as global memory, where voting and synchronization take place so some access-time burden is introduced; nevertheless, the speed of the global memory is much faster than disk access, so this level is used for page swapping with local memory to keep the most-used data in the fastest area, rather than employing disk for the first level of demand paging.

One of the features of the disclosed embodiment of the invention is ability to replace faulty components, such as CPU modules or memory modules, without shutting down the system. Thus, the system is available for continuous use even though components may fail and have to be replaced. In addition, the ability to obtain a high level of fault tolerance with fewer system components, e.g., no fault-tolerant clocking needed, only two memory modules needed instead of three, voting circuits minimized, etc., means that there are fewer components to fail, and so the reliability is enhanced. That is, there are fewer failures because there are fewer components, and when there are failures the components are isolated to allow the system to keep running, while the components can be replaced without system shut-down.

The CPUs of this system preferably use a commercially-available high-performance microprocessor chip for which operating systems such as Unix™ are available. The parts of the system which make it fault-tolerant are either transparent to the operating system or easily adapted to the operating system. Accordingly, a high-performance fault-tolerant system is provided which allows comparability with contemporary widely-used multi-tasking operating system and applications software.

According to one embodiment, the present invention is directed to a method and apparatus for loosely synchronizing a plurality of processors. The apparatus according to the invention allows two or more processors to be configured in a fault detecting or fault tolerant arrangement without the need for a fault tolerant clock circuit. The processors are free to execute the same algorithm at different

speeds, whether due to differences in clock rates or the occurrence of "extra" clock cycles. "Extra" clock cycles may occur because of error retries, variations in cache hit rates, or as a result of asynchronous logic, and they represent clock cycles which ordinarily do not occur when executing a particular program. External interrupts are synchronized in a manner such that each processor responds to the interrupt at the same point in its execution with due regard for maximum interrupt latency time.

In one embodiment of the present invention, each processor runs off of its own independent clock. The processor indicates the occurrence of a prescribed processor event on one line and receives signals on another line for initiating a wait state. A processor event may be defined either explicitly or implicitly by the code running in the processor, and it is preferable to generate one processor event signal for each microprocessor write operation. Each processor has a counter, termed an event counter, which counts the number of processor events indicated since the last time the processors were synchronized.

In this embodiment, the processors typically are synchronized whenever an external interrupt occurs, although the system designer is free to define any synchronizing event. When an event requiring synchronization is detected by a sync logic circuit associated with the processor, the sync logic circuit generates a wait signal after the next processor event. A compare circuit associated with each processor tests the other event counters in the system and determines whether its associated processor is behind the others. If so, the sync logic circuit removes the wait signal until the next processor event. The compare circuit then rechecks to see if its associated processor is still behind. The processor is finally stopped when its event counter matches the event counter for the fastest processor. The process continues until all processors are stopped with each event counter having the same value. When this point is reached, the processors are then all stopped at the same point in the program. The wait signal is removed, the interrupt line to each processor is asserted, and all processors are restarted to handle the synchronizing event.

If no synchronizing event occurs before an event counter reaches its maximum value, an overflow of the event counter forces resynchronization. The affected processor waits until the other processors event counters also overflow before continuing. On the other hand, if a synchronizing event occurs but the processor events do not occur often enough to satisfy worst-case interrupt latency times, another counter, termed a cycle counter, is provided for counting the number of clock cycles

since the last processor event. The cycle counter is set to overflow at a point before maximum interrupt latency time is exceeded. An overflow of the cycle counter forces resynchronization by generating an internal synchronization request signal and an interrupt signal. When the processor services the interrupt, the code within the interrupt routine forces an event to be generated. The internally generated synchronization request signal thus causes resynchronization to the event generated by the interrupt routine. The processors then may serve the pending interrupt.

In another embodiment, "run" cycles of a CPU are counted in an event counter, which is in this case a cycle counter; that is, all non-stall cycles (where the pipeline advances) are the events being counted. Then, upon a synchronization request in the form of an external interrupt, the CPUs are kept in synchronization by waiting until each CPU is at the same event (executing the instruction at the same cycle count) before the interrupt is presented to that CPU. Thus, a CPU may receive the interrupt at a different "real" time, but it will receive the interrupt at the same time as the others measured in what instruction is being executed, so CPUs are not necessarily brought back into "real time" synchronization by the interrupt. This interrupt sync method is used along with another synchronization technique, in that external memory references are voted and the memory reference is not implemented until all CPUs have made the same reference (or a failure is detected), thus forcing real-time synchronization. In addition, overflow of the cycle counter causes synchronization, so that if a memory reference does not occur within some selected period (represented by the length of the cycle count register) then a synchronization operation will be performed to keep the CPUs from drifting too far apart.

BRIEF DESCRIPTION OF THE DRAWINGS

The features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as other features and advantages thereof, may best be understood by reference to the detailed description of a specific embodiment which follows, when read in conjunction with the accompanying drawings, wherein:

Figure 1 is an electrical diagram in block form of a computer system according to one embodiment of the invention;

Figure 2 is an electrical schematic diagram in block form of one of the CPUs of the system of Figure 1;

Figure 3 is an electrical schematic diagram in block form of one of the microprocessor chip used in the CPU of Figure 2;

Figures 4 and 5 are timing diagrams showing events occurring in the CPU of Figures 2 and 3 as a function of time;

Figure 6 is an electrical schematic diagram in block form of one of the memory modules in the computer system of Figure 1;

Figure 7 is a timing diagram showing events occurring on the CPU to memory busses in the system of Figure 1;

Figure 8 is an electrical schematic diagram in block form of one of the I/O processors in the computer system of Figure 1;

Figure 9 is a timing diagram showing events vs. time for the transfer protocol between a memory module and an I/O processor in the system of Figure 1;

Figure 10 is a timing diagram showing events vs. time for execution of instructions in the CPUs of Figures 1, 2 and 3;

Figure 10a is a detail view of a part of the diagram of Figure 10;

Figures 11 and 12 are timing diagrams similar to Figure 10 showing events vs. time for execution of instructions in the CPUs of Figures 1, 2 and 3;

Figure 13 is an electrical schematic diagram in block form of the interrupt synchronization circuit used in the CPU of Figure 2;

Figures 14, 15, 16 and 17 are timing diagrams like Figures 10 or 11 showing events vs. time for execution of instructions in the CPUs of Figures 1, 2 and 3 when an interrupt occurs, illustrating various scenarios;

Figure 18 is a physical memory map of the memories used in the system of Figures 1, 2, 3 and 6;

Figure 19 is a virtual memory map of the CPUs used in the system of Figures 1, 2, 3 and 6;

Figure 20 is a diagram of the format of the virtual address and the TLB entries in the micro-processor chips in the CPU according to Figure 2 or 3;

Figure 21 is an illustration of the private memory locations in the memory map of the global memory modules in the system of Figures 1, 2, 3 and 6;

Figure 22 is an electrical diagram of a fault-tolerant power supply used with the system of the invention according to one embodiment;

Fig. 23 is a conceptual block diagram of an embodiment of a data processing system according to another example of an embodiment of the present invention;

Fig. 24 is a diagram of a processing sequence of the data processing system illustrated in Fig. 23;

Fig. 25 is a conceptual block diagram of an embodiment of a CPU illustrated in Fig. 23;

Fig. 26A-26C are diagrams illustrating processor synchronizing procedures according to the embodiment of Figure 23;

Fig. 27 is a flow chart illustrating processor synchronization according to the embodiment of Figure 23; and

Figure 28 is an electrical diagram (like Figure 1) of a system according to another embodiment of the invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENT

With reference to Figure 1, a computer system using features of the invention is shown in one embodiment having three identical processors 11, 12 and 13, referred to as CPU-A, CPU-B and CPU-C, which operate as one logical processor, all three typically executing the same instruction stream; the only time the three processors are not executing the same instruction stream is in such operations as power-up self test, diagnostics and the like. The three processors are coupled to two memory modules 14 and 15, referred to as Memory-# 1 and Memory-#2, each memory storing the same data in the same address space. In a preferred embodiment, each one of the processors 11, 12 and 13 contains its own local memory 16, as well, accessible only by the processor containing this memory.

Each one of the processors 11, 12 and 13, as well as each one of the memory modules 14 and 15, has its own separate clock oscillator 17; in this embodiment, the processors are not run in "lock step", but instead are loosely synchronized by a method such as is set forth in the above-mentioned application Ser. No. 118,503, i.e., using events such as external memory references to bring the CPUs into synchronization. External interrupts are synchronized among the three CPUs by a technique employing a set of busses 18 for coupling the interrupt requests and status from each of the processors to the other two; each one of the processors CPU-A, CPU-B and CPU-C is responsive to the three interrupt requests, its own and the two received from the other CPUs, to present an interrupt to the CPUs at the same point in the execution stream. The memory modules 14 and 15 vote the memory references, and allow a memory reference to proceed only when all three CPUs have made the same request (with provision for faults). In this manner, the processors are synchronized at the time of external events (memory references), resulting in the processors typically executing the same instruction stream, in the same sequence, but not necessarily during aligned clock cycles in the time between synchronization events. In addition, external interrupts are synchronized to be

executed at the same point in the instruction stream of each CPU.

The CPU-A processor 11 is connected to the Memory-#1 module 14 and to the Memory-#2 module 15 by a bus 21; likewise the CPU-B is connected to the modules 14 and 15 by a bus 22, and the CPU-C is connected to the memory modules by a bus 23. These busses 21, 22, 23 each include a 32-bit multiplexed address/data bus, a command bus, and control lines for address and data strobes. The CPUs have control of these busses 21, 22 and 23, so there is no arbitration, or bus-request and bus-grant.

Each one of the memory modules 14 and 15 is separately coupled to a respective input/output bus 24 or 25, and each of these busses is coupled to two (or more) input/output processors 26 and 27. The system can have multiple I/O processors as needed to accommodate the I/O devices needed for the particular system configuration. Each one of the input/output processors 26 and 27 is connected to a bus 28, which may be of a standard configuration such as a VMEbus™, and each bus 28 is connected to one or more bus interface modules 29 for interface with a standard I/O controller 30. Each bus interface module 29 is connected to two of the busses 28, so failure of one I/O processor 26 or 27, or failure of one of the bus channels 28, can be tolerated. The I/O processors 26 and 27 can be addressed by the CPUs 11, 12 and 13 through the memory modules 14 and 15, and can signal an interrupt to the CPUs via the memory modules. Disk drives, terminals with CRT screens and keyboards, and network adapters, are typical peripheral devices operated by the controllers 30. The controllers 30 may make DMA-type references to the memory modules 14 and 15 to transfer blocks of data. Each one of the I/O processors 26, 27, etc., has certain individual lines directly connected to each one of the memory modules for bus request, bus grant, etc.; these point-to-point connections are called "radials" and are included in a group of radial lines 31.

A system status bus 32 is individually connected to each one of the CPUs 11, 12 and 13, to each memory module 14 and 15, and to each of the I/O processors 26 and 27, for the purpose of providing information on the status of each element. This status bus provides information about which of the CPUs, memory modules and I/O processors is currently in the system and operating properly.

An acknowledge/status bus 33 connecting the three CPUs and two memory modules includes individual lines by which the modules 14 and 15 send acknowledge signals to the CPUs when memory requests are made by the CPUs, and at the same time a status field is sent to report on the

status of the command and whether it executed correctly. The memory modules not only check parity on data read from or written to the global memory, but also check parity on data passing through the memory modules to or from the I/O busses 24 and 25, as well as checking the validity of commands. It is through the status lines in bus 33 that these checks are reported to the CPUs 11, 12 and 13, so if errors occur a fault routine can be entered to isolate a faulty component.

Even though both memory modules 14 and 15 are storing the same data in global memory, and operating to perform every memory reference in duplicate, one of these memory modules is designated as primary and the other as back-up, at any given time. Memory write operations are executed by both memory modules so both are kept current, and also a memory read operation is executed by both, but only the primary module actually loads the read-data back onto the busses 21, 22 and 23, and only the primary memory module controls the arbitration for multi-master busses 24 and 25. To keep the primary and back-up modules executing the same operations, a bus 34 conveys control information from primary to back-up. Either module can assume the role of primary at boot-up, and the roles can switch during operation under software control; the roles can also switch when selected error conditions are detected by the CPUs or other error-responsive parts of the system.

Certain interrupts generated in the CPUs are also voted by the memory modules 14 and 15. When the CPUs encounter such an interrupt condition (and are not stalled), they signal an interrupt request to the memory modules by individual lines in an interrupt bus 35, so the three interrupt requests from the three CPUs can be voted. When all interrupts have been voted, the memory modules each send a voted-interrupt signal to the three CPUs via bus 35. This voting of interrupts also functions to check on the operation of the CPUs. The three CPUs synch the voted interrupt CPU interrupt signal via the inter-CPU bus 18 and present the interrupt to the processors at a common point in the instruction stream. This interrupt synchronization is accomplished without stalling any of the CPUs.

CPU Module:

Referring now to Figure 2, one of the processors 11, 12 or 13 is shown in more detail. All three CPU modules are of the same construction in a preferred embodiment, so only CPU-A will be described here. In order to keep costs within a competitive range, and to provide ready access to already-developed software and operating systems, it is preferred to use a commercially-available

microprocessor chip, and any one of a number of devices may be chosen. The RISC (reduced instruction set) architecture has some advantage in implementing the loose synchronization as will be described, but more-conventional CISC (complex instruction set) microprocessors such as Motorola 68030 devices or Intel 80386 devices (available in 20-MHz and 25-MHz speeds) could be used. High-speed 32-bit RISC microprocessor devices are available from several sources in three basic types; Motorola produces a device as part number 88000, MIPS Computer Systems, Inc. and others produce a chip set referred to as the MIPS type, and Sun Microsystems has announced a so-called SPARCTM type (scalable processor architecture). Cypress Semiconductor of San Jose, California, for example, manufactures a microprocessor referred to as part number CY7C601 providing 20-MIPS (million instructions per second), clocked at 33-MHz, supporting the SPARC standard, and Fujitsu manufactures a CMOS RISC microprocessor, part number S-25, also supporting the SPARC standard.

The CPU board or module in the illustrative embodiment, used as an example, employs a microprocessor chip 40 which is in this case an R2000 device designed by MIPS Computer Systems, Inc., and also manufactured by Integrated Device Technology, Inc. The R2000 device is a 32-bit processor using RISC architecture to provide high performance, e.g., 12-MIPS at 16.67-MHz clock rate. Higher-speed versions of this device may be used instead, such as the R3000 that provides 20-MIPS at 25-MHz clock rate. The processor 40 also has a co-processor used for memory management, including a translation lookaside buffer to cache translations of logical to physical addresses. The processor 40 is coupled to a local bus having a data bus 41, an address bus 42 and a control bus 43. Separate instruction and data cache memories 44 and 45 are coupled to this local bus. These caches are each of 64K-byte size, for example, and are accessed within a single clock cycle of the processor 40. A numeric or floating point co-processor 46 is coupled to the local bus if additional performance is needed for these types of calculations; this numeric processor device is also commercially available from MIPS Computer Systems as part number R2010. The local bus 41, 42, 43, is coupled to an internal bus structure through a write buffer 50 and a read buffer 51. The write buffer is a commercially available device, part number R2020, and functions to allow the processor 40 to continue to execute Run cycles after storing data and address in the write buffer 50 for a write operation, rather than having to execute stall cycles while the write is completing.

In addition to the path through the write buffer 50, a path is provided to allow the processor 40 to

execute write operations bypassing the write buffer 50. This path is a write buffer bypass 52 allows the processor, under software selection, to perform synchronous writes. If the write buffer bypass 52 is enabled (write buffer 50 not enabled) and the processor executes a write then the processor will stall until the write completes. In contrast, when writes are executed with the write buffer bypass 52 disabled the processor will not stall because data is written into the write buffer 50 (unless the write buffer is full). If the write buffer 50 is enabled when the processor 40 performs a write operation, the write buffer 50 captures the output data from bus 41 and the address from bus 42, as well as controls from bus 43. The write buffer 50 can hold up to four such data-address sets while it waits to pass the data on to the main memory. The write buffer runs synchronously with the clock 17 of the processor chip 40, so the processor-to-buffer transfers are synchronous and at the machine cycle rate of the processor. The write buffer 50 signals the processor if it is full and unable to accept data. Read operations by the processor 40 are checked against the addresses contained in the four-deep write buffer 50, so if a read is attempted to one of the data words waiting in the write buffer to be written to memory 16 or to global memory, the read is stalled until the write is completed.

The write and read buffers 50 and 51 are coupled to an internal bus structure having a data bus 53, an address bus 54 and a control bus 55. The local memory 16 is accessed by this internal bus, and a bus interface 56 coupled to the internal bus is used to access the system bus 21 (or bus 22 or 23 for the other CPUs). The separate data and address busses 53 and 54 of the internal bus (as derived from busses 41 and 42 of the local bus) are converted to a multiplexed address/data bus 57 in the system bus 21, and the command and control lines are correspondingly converted to command lines 58 and control lines 59 in this external bus.

The bus interface unit 56 also receives the acknowledge/status lines 33 from the memory modules 14 and 15. In these lines 33, separate status lines 33-1 or 33-2 are coupled from each of the modules 14 and 15, so the responses from both memory modules can be evaluated upon the event of a transfer (read or write) between CPUs and global memory, as will be explained.

The local memory 16, in one embodiment, comprises about 8-Mbyte of RAM which can be accessed in about three or four of the machine cycles of processor 40, and this access is synchronous with the clock 17 of this CPU, whereas the memory access time to the modules 14 and 15 is much greater than that to local memory, and this access to the memory modules 14 and 15 is asyn-

chronous and subject to the synchronization overhead imposed by waiting for all CPUs to make the request then voting. For comparison, access to a typical commercially-available disk memory through the I/O processors 26, 27 and 29 is measured in milliseconds, i.e., considerably slower than access to the modules 14 and 15. Thus, there is a hierarchy of memory access by the CPU chip 40, the highest being the instruction and data caches 44 and 45 which will provide a hit ratio of perhaps 95% when using 64-KByte cache size and suitable fill algorithms. The second highest is the local memory 16, and again by employing contemporary virtual memory management algorithms a hit ratio of perhaps 95% is obtained for memory references for which a cache miss occurs but a hit in local memory 16 is found, in an example where the size of the local memory is about 8-MByte. The net result, from the standpoint of the processor chip 40, is that perhaps greater than 99% of memory references (but not I/O references) will be synchronous and will occur in either the same machine cycle or in three or four machine cycles.

The local memory 16 is accessed from the internal bus by a memory controller 60 which receives the addresses from address bus 54, and the address strobes from the control bus 55, and generates separate row and column addresses, and RAS and CAS controls, for example, if the local memory 16 employs DRAMs with multiplexed addressing, as is usually the case. Data is written to or read from the local memory via data bus 53. In addition, several local registers 61, as well as non-volatile memory 62 such as NVRAMs, and high-speed PROMs 63, as may be used by the operating system, are accessed by the internal bus; some of this part of the memory is used only at power-on, some is used by the operating system and may be almost continuously within the cache 44, and other may be within the non-cached part of the memory map.

External interrupts are applied to the processor 40 by one of the pins of the control bus 43 or 55 from an interrupt circuit 65 in the CPU module of Figure 2. This type of interrupt is voted in the circuit 65, so that before an interrupt is executed by the processor 40 it is determined whether or not all three CPUs are presented with the interrupt; to this end, the circuit 65 receives interrupt pending inputs 66 from the other two CPUs 12 and 13, and sends an interrupt pending signal to the other two CPUs via line 67, these lines being part of the bus 18 connecting the three CPUs 11, 12 and 13 together. Also, for voting other types of interrupts, specifically CPU-generated interrupts, the circuit 65 can send an interrupt request from this CPU to both of the memory modules 14 and 15 by a line 68 in the bus 35, then receive separate voted-

interrupt signals from the memory modules via lines 69 and 70; both memory modules will present the external interrupt to be acted upon. An interrupt generated in some external source such as a keyboard or disk drive on one of the I/O channels 28, for example, will not be presented to the interrupt pin of the chip 40 from the circuit 65 until each one of the CPUs 11, 12 and 13 is at the same point in the instruction stream, as will be explained.

Since the processors 40 are clocked by separate clock oscillators 17, there must be some mechanism for periodically bringing the processors 40 back into synchronization. Even though the clock oscillators 17 are of the same nominal frequency, e.g., 16.67-MHz, and the tolerance for these devices is about 25-ppm (parts per million), the processors can potentially become many cycles out of phase unless periodically brought back into synch. Of course, every time an external interrupt occurs the CPUs will be brought into synch in the sense of being interrupted at the same point in their instruction stream (due to the interrupt synch mechanism), but this does not help bring the cycle count into synch. The mechanism of voting memory references in the memory modules 14 and 15 will bring the CPUs into synch (in real time), as will be explained. However, some conditions result in long periods where no memory reference occurs, and so an additional mechanism is used to introduce stall cycles to bring the processors 40 back into synch. A cycle counter 71 is coupled to the clock 17 and the control pins of the processor 40 via control bus 43 to count machine cycles which are Run cycles (but not Stall cycles). This counter 71 includes a count register having a maximum count value selected to represent the period during which the maximum allowable drift between CPUs would occur (taking into account the specified tolerance for the crystal oscillators); when this count register overflows action is initiated to stall the faster processors until the slower processor or processors catch up. This counter 71 is reset whenever a synchronization is done by a memory reference to the memory modules 14 and 15. Also, a refresh counter 72 is employed to perform refresh cycles on the local memory 16, as will be explained. In addition, a counter 73 counts machine cycle which are Run cycles but not Stall cycles, like the counter 71 does, but this counter 73 is not reset by a memory reference; the counter 73 is used for interrupt synchronization as explained below, and to this end produces the output signals CC-4 and CC-8 to the interrupt synchronization circuit 65.

The processor 40 has a RISC instruction set which does not support memory-to-memory instructions, but instead only memory-to-register or register-to-memory instructions (i.e., load or store).

It is important to keep frequently-used data and the currently-executing code in local memory. Accordingly, a block-transfer operation is provided by a DMA state machine 74 coupled to the bus interface 56. The processor 40 writes a word to a register in the DMA circuit 74 to function as a command, and writes the starting address and length of the block to registers in this circuit 74. In one embodiment, the microprocessor stalls while the DMA circuit takes over and executes the block transfer, producing the necessary addresses, commands and strobes on the busses 53-55 and 21. The command executed by the processor 40 to initiate this block transfer can be a read from a register in the DMA circuit 74. Since memory management in the Unix operating system relies upon demand paging, these block transfers will most often be pages being moved between global and local memory and I/O traffic. A page is 4-KBytes. Of course, the busses 21, 22 and 23 support single-word read and write transfers between CPUs and global memory; the block transfers referred to are only possible between local and global memory.

The Processor:

Referring now to Figure 3, the R2000 or R3000 type of microprocessor 40 of the example embodiment is shown in more detail. This device includes a main 32-bit CPU 75 containing thirty-two 32-bit general purpose registers 76, a 32-bit ALU 77, a zero-to-64 bit shifter 78, and a 32-by-32 multiply/divide circuit 79. This CPU also has a program counter 80 along with associated incrementer and adder. These components are coupled to a processor bus structure 81, which is coupled to the local data bus 41 and to an instruction decoder 82 with associated control logic to execute instructions fetched via data bus 41. The 32-bit local address bus 42 is driven by a virtual memory management arrangement including a translation lookaside buffer (TLB) 83 within an on-chip memory-management coprocessor. The TLB 83 contains sixty-four entries to be compared with a virtual address received from the microprocessor block 75 via virtual address bus 84. The low-order 16-bit part 85 of the bus 42 is driven by the low-order part of this virtual address bus 84, and the high-order part is from the bus 84 if the virtual address is used as the physical address, or is the tag entry from the TLB 83 via output 86 if virtual addressing is used and a hit occurs. The control lines 43 of the local bus are connected to pipeline and bus control circuitry 87, driven from the internal bus structure 81 and the control logic 82.

The microprocessor block 75 in the processor 40 is of the RISC type in that most instructions execute in one machine cycle, and the instruction

set uses register-to-register and load/store instructions rather than having complex instructions involving memory references along with ALU operations. There are no complex addressing schemes included as part of the instruction set, such as "add the operand whose address is the sum of the contents of register A1 and register A2 to the operand whose address is found at the main memory location addressed by the contents of register B, and store the result in main memory at the location whose address is found in register C." Instead, this operation is done in a number of simple register-to-register and load/store instructions: add register A2 to register A1; load register B1 from memory location whose address is in register B; add register A1 and register B1; store register B1 to memory location addressed by register C. Optimizing compiler techniques are used to maximize the use of the thirty-two registers 76, i.e., assure that most operations will find the operands already in the register set. The load instructions actually take longer than one machine cycle, and to account for this a latency of one instruction is introduced; the data fetched by the load instruction is not used until the second cycle, and the intervening cycle is used for some other instruction, if possible.

The main CPU 75 is highly pipelined to facilitate the goal of averaging one instruction execution per machine cycle. Referring to Figure 4, a single instruction is executed over a period including five machine cycles, where a machine cycle is one clock period or 60-nsec for a 16.67-MHz clock 17. These five cycles or pipe stages are referred to as IF (instruction fetch from I-cache 44), RD (read operands from register set 76), ALU (perform the required operation in ALU 77), MEM (access D-cache 45 if required), and WB (write back ALU result to register file 76). As seen in Figure 5, these five pipe stages are overlapped so that in a given machine cycle, cycle-5 for example, instruction I#5 is in its first or IF pipe stage and instruction I#1 is in its last or WB stage, while the other instructions are in the intervening pipe stages.

Memory Module:

With reference to Figure 6, one of the memory modules 14 or 15 is shown in detail. Both memory modules are of the same construction in a preferred embodiment, so only the Memory# 1 module is shown. The memory module includes three input/output ports 91, 92 and 93 coupled to the three busses 21, 22 and 23 coming from the CPUs 11, 12 and 13, respectively. Inputs to these ports are latched into registers 94, 95 and 96 each of which has separate sections to store data, address, command and strobes for a write operation, or

address, command and strobes for a read operation. The contents of these three registers are voted by a vote circuit 100 having inputs connected to all sections of all three registers. If all three of the CPUs 11, 12 and 13 make the same memory request (same address, same command), as should be the case since the CPUs are typically executing the same instruction stream, then the memory request is allowed to complete; however, as soon as the first memory request is latched into any one of the three latches 94, 95 or 96, it is passed on immediately to begin the memory access. To this end, the address, data and command are applied to an internal bus including data bus 101, address bus 102 and control bus 103. From this internal bus the memory request accesses various resources, depending upon the address, and depending upon the system configuration.

In one embodiment, a large DRAM 104 is accessed by the internal bus, using a memory controller 105 which accepts the address from address bus 102 and memory request and strobes from control bus 103 to generate multiplexed row and column addresses for the DRAM so that data input/output is provided on the data bus 101. This DRAM 104 is also referred to as global memory, and is of a size of perhaps 32-MByte in one embodiment. In addition, the internal bus 101-103 can access control and status registers 106, a quantity of non-volatile RAM 107, and write-protect RAM 108. The memory reference by the CPUs can also bypass the memory in the memory module 14 or 15 and access the I/O busses 24 and 25 by a bus interface 109 which has inputs connected to the internal bus 101-103. If the memory module is the primary memory module, a bus arbitrator 110 in each memory module controls the bus interface 109. If a memory module is the backup module, the bus 34 controls the bus interface 109.

A memory access to the DRAM 104 is initiated as soon as the first request is latched into one of the latches 94, 95 or 96, but is not allowed to complete unless the vote circuit 100 determines that a plurality of the requests are the same, with provision for faults. The arrival of the first of the three requests causes the access to the DRAM 104 to begin. For a read, the DRAM 104 is addressed, the sense amplifiers are strobed, and the data output is produced at the DRAM outputs, so if the vote is good after the third request is received then the requested data is ready for immediate transfer back to the CPUs. In this manner, voting is overlapped with DRAM access.

Referring to Figure 7, the busses 21, 22 and 23 apply memory requests to ports 91, 92 and 93 of the memory modules 14 and 15 in the format illustrated. Each of these busses consists of thirty-two bidirectional multiplexed address/data lines,

thirteen unidirectional command lines, and two strobes. The command lines include a field which specifies the type of bus activity, such as read, write, block transfer, single transfer, I/O read or write, etc. Also, a field functions as a byte enable for the four bytes. The strobes are AS, address strobe, and DS, data strobe. The CPUs 11, 12 and 13 each control their own bus 21, 22 or 23; in this embodiment, these are not multi-master busses, there is no contention or arbitration. For a write, the CPU drives the address and command onto the bus in one cycle along with the address strobe AS (active low), then in a subsequent cycle (possibly the next cycle, but not necessarily) drives the data onto the address/data lines of the bus at the same time as a data strobe DS. The address strobe AS from each CPU causes the address and command then appearing at the ports 91, 92 or 93 to be latched into the address and command sections of the registers 94, 95 and 96, as these strobes appear, then the data strobe DS causes the data to be latched. When a plurality (two out of three in this embodiment) of the busses 21, 22 and 23 drive the same memory request into the latches 94, 95 and 96, the vote circuit 100 passes on the final command to the bus 103 and the memory access will be executed; if the command is a write, an acknowledge ACK signal is sent back to each CPU by a line 112 (specifically line 112-1 for Memory#1 and line 112-2 for Memory#2) as soon as the write has been executed, and at the same time status bits are driven via acknowledge/status bus 33 (specifically lines 33-1 for Memory#1 and lines 33-2 for Memory#2) to each CPU at time T3 of Figure 7. The delay T4 between the last strobe DS (or AS if a read) and the ACK at T3 is variable, depending upon how many cycles out of synch the CPUs are at the time of the memory request, and depending upon the delay in the voting circuit and the phase of the internal independent clock 17 of the memory module 14 or 15 compared to the CPU clocks 17. If the memory request issued by the CPUs is a read, then the ACK signal on lines 112-1 and 112-2 and the status bits on lines 33-1 and 33-2 will be sent at the same time as the data is driven to the address/data bus, during time T3; this will release the stall in the CPUs and thus synchronize the CPU chips 40 on the same instruction. That is, the fastest CPU will have executed more stall cycles as it waited for the slower ones to catch up, then all three will be released at the same time, although the clocks 17 will probably be out of phase; the first instruction executed by all three CPUs when they come out of stall will be the same instruction.

All data being sent from the memory module 14 or 15 to the CPUs 11, 12 and 13, whether the data is read data from the DRAM 104 or from the

memory locations 106-108, or is I/O data from the busses 24 and 25, goes through a register 114. This register is loaded from the internal data bus 101, and an output 115 from this register is applied to the address/data lines for busses 21, 22 and 23 at ports 91, 92 and 93 at time T3. Parity is checked when the data is loaded to this register 114. All data written to the DRAM 104, and all data on the I/O busses, has parity bits associated with it, but the parity bits are not transferred on busses 21, 22 and 23 to the CPU modules. Parity errors detected at the read register 114 are reported to the CPU via the status busses 33-1 and 33-2. Only the memory module 14 or 15 designated as primary will drive the data in its register 114 onto the busses 21, 22 and 23. The memory module designated as back-up or secondary will complete a read operation all the way up to the point of loading the register 114 and checking parity, and will report status on buses 33-1 and 33-2, but no data will be driven to the busses 21, 22 and 23.

A controller 117 in each memory module 14 or 15 operates as a state machine clocked by the clock oscillator 17 for this module and receiving the various command lines from bus 103 and busses 21-23, etc., to generate control bits to load registers and busses, generate external control signals, and the like. This controller also is connected to the bus 34 between the memory modules 14 and 15 which transfers status and control information between the two. The controller 117 in the module 14 or 15 currently designated as primary will arbitrate via arbitrator 110 between the I/O side (interface 109) and the CPU side (ports 91-93) for access to the common bus 101-103. This decision made by the controller 117 in the primary memory module 14 or 15 is communicated to the controller 117 of other memory module by the lines 34, and forces the other memory module to execute the same access.

The controller 117 in each memory module also introduces refresh cycles for the DRAM 104, based upon a refresh counter 118 receiving pulses from the clock oscillator 17 for this module. The DRAM must receive 512 refresh cycles every 8-msec, so on average there must be a refresh cycle introduced about every 15-microsec. The counter 118 thus produces an overflow signal to the controller 117 every 15-microsec., and if an idle condition exists (no CPU access or I/O access executing) a refresh cycle is implemented by a command applied to the bus 103. If an operation is in progress, the refresh is executed when the current operation is finished. For lengthy operations such as block transfers used in memory paging, several refresh cycles may be backed up and execute in a burst mode after the transfer is completed; to this end, the number of overflows -of counter 118 since

the last refresh cycle are accumulated in a register associated with the counter 118.

Interrupt requests for CPU-generated interrupts are received from each CPU 11, 12 and 13 individually by lines 68 in the interrupt bus 35; these interrupt requests are sent to each memory module 14 and 15. These interrupt request lines 68 in bus 35 are applied to an interrupt vote circuit 119 which compares the three requests and produces a voted interrupt signal on outgoing line 69 of the bus 35. The CPUs each receive a voted interrupt signal on the two lines 69 and 70 (one from each module 14 and 15) via the bus 35. The voted interrupts from each memory module 14 and 15 are ORed and presented to the interrupt synchronizing circuit 65. The CPUs, under software control, decide which interrupts to service. External interrupts, generated in the I/O processors or I/O controllers, are also signalled to the CPUs through the memory modules 14 and 15 via lines 69 and 70 in bus 35, and likewise the CPUs only respond to an interrupt from the primary module 14 or 15.

I/O Processor:

Referring now to Figure 8, one of the I/O processors 26 or 27 is shown in detail. The I/O processor has two identical ports, one port 121 to the I/O bus 24 and the other port 122 to the I/O bus 25. Each one of the I/O busses 24 and 25 consists of: a 36-bit bidirectional multiplexed address/data bus 123 (containing 32-bits plus 4-bits parity), a bidirectional command bus 124 defining the read, write, block read, block write, etc., type of operation that is being executed, an address line that designates which location is being addressed, either internal to I/O processor or on busses 28, and the byte mask, and finally control lines 125 including address strobe, data strobe, address acknowledge and data acknowledge. The radial lines in bus 31 include individual lines from each I/O processor to each memory module: bus request from I/O processor to the memory modules, bus grant from the memory modules to the I/O processor, interrupt request lines from I/O processor to memory module, and a reset line from memory to I/O processor. Lines to indicate which memory module is primary are connected to each I/O processor via the system status bus 32. A controller or state machine 126 in the I/O processor of Figure 8 receives the command, control, status and radial lines and internal data, and command lines from the busses 28, and defines the internal operation of the I/O processor, including operation of latches 127 and 128 which receive the contents of busses 24 and 25 and also hold information for transmitting onto the busses.

Transfer on the busses 24 and 25 from memory module to I/O processor uses a protocol as

shown in Figure 9 with the address and data separately acknowledged. The arbitrator circuit 110 in the memory module which is designated primary performs the arbitration for ownership of the I/O busses 24 and 25. When a transfer from CPUs to I/O is needed, the CPU request is presented to the arbitration logic 110 in the memory module. When the arbiter 110 grants this request the memory modules apply the address and command to busses 123 and 124 (of both busses 24 and 25) at the same time the address strobe is asserted on bus 125 (of both busses 24 and 25) in time T1 of Figure 9; when the controller 126 has caused the address to be latched into latches 127 or 128, the address acknowledge is asserted on bus 125, then the memory modules place the data (via both busses 24 and 25) on the bus 123 and a data strobe on lines 125 in time T2, following which the controller causes the data to be latched into both latches 127 and 128 and a data acknowledge signal is placed upon the lines 125, so upon receipt of the data acknowledge, both of the memory modules release the bus 24, 25 by de-asserting the address strobe signal. The I/O processor then deasserts the address acknowledge signal.

For transfers from I/O processor to the memory module, when the I/O processor needs to use the I/O bus, it asserts a bus request by a line in the radial bus 31, to both busses 24 and 25, then waits for a bus grant signal from an arbitrator circuit 110 in the primary memory module 14 or 15, the bus grant line also being one of the radials. When the bus grant has been asserted, the controller 126 then waits until the address strobe and address acknowledge signals on busses 125 are deasserted (i.e., false) meaning the previous transfer is completed. At that time, the controller 126 causes the address to be applied from latches 127 and 128 to lines 123 of both busses 24 and 25, the command to be applied to lines 124, and the address strobe to be applied to the bus 125 of both busses 24 and 25. When address acknowledge is received from both busses 24 and 25, these are followed by applying the data to the address/data busses, along with data strobes, and the transfer is completed with a data acknowledge signals from the memory modules to the I/O processor.

The latches 127 and 128 are coupled to an internal bus 129 including an address bus 129a, and data bus 129b and a control bus 129c, which can address internal status and control registers 130 used to set up the commands to be executed by the controller state machine 126, to hold the status distributed by the bus 32, etc. These registers 130 are addressable for read or write from the CPUs in the address space of the CPUs. A bus interface 131 communicates with the VMEbus 28, under control of the controller 126. The bus 28

includes an address bus 28a, a data bus 28b, a control bus 28c, and radials 28d, and all of these lines are communicated through the bus interface modules 29 to the I/O controllers 30; the bus interface module 29 contains a multiplexer 132 to allow only one set of bus lines 28 (from one I/O processor or the other but not both) drive the controller 30. Internal to the controller 30 are command, control, status and data registers 133 which (as is standard practice for peripheral controllers of this type) are addressable from the CPUs 11, 12 and 13 for read and write to initiate and control operations in I/O devices.

Each one of the I/O controllers 30 on the VMEbuses 28 has connections via a multiplexer 132 in the BIM 29 to both I/O processors 26 and 27 and can be controlled by either one, but is bound to one or the other by the program executing in the CPUs. A particular address (or set of addresses) is established for control and data-transfer registers 133 representing each controller 30, and these addresses are maintained in an I/O page table (normally in the kernel data section of local memory) by the operating system. These addresses associate each controller 30 as being accessible only through either I/O processor #1 or #2, but not both. That is, a different address is used to reach a particular register 133 via I/O processor 26 compared to I/O processor 27. The bus interface 131 (and controller 126) can switch the multiplexer 132 to accept bus 28 from one or the other, and this is done by a write to the registers 130 of the I/O processors from the CPUs. Thus, when the device driver is called up to access this controller 30, the operating system uses these addresses in the page table to do it. The processors 40 access the controllers 30 by I/O writes to the control and data-transfer registers 133 in these controllers using the write buffer bypass path 52, rather than through the write buffer 50, so these are synchronous writes, voted by circuits 100, passed through the memory modules to the busses 24 or 25, thus to the selected bus 28; the processors 40 stall until the write is completed. The I/O processor board of Figure 8 is configured to detect certain failures, such as improper commands, time-outs where no response is received over VMEbus 28, parity-checked data if implemented, etc., and when one of these failures is detected the I/O processor quits responding to bus traffic, i.e., quits sending address acknowledge and data acknowledge as discussed above with reference to Figure 9. This is detected by the bus interface 56 as a bus fault, resulting in an interrupt as will be explained, and self-correcting action if possible.

Error Recovery:

The sequence used by the CPUs 11, 12 and 13 to evaluate responses by the memory modules 14 and 15 to transfers via busses 21, 22 and 23 will now be described. This sequence is defined by the state machine in the bus interface units 56 and in code executed by the CPUs.

In case one, for a read transfer, it is assumed that no data errors are indicated in the status bits on lines 33 from the primary memory. Here, the stall begun by the memory reference is ended by asserting a Ready signal via control bus 55 and 43 to allow instruction execution to continue in each microprocessor 40. But, another transfer is not started until acknowledge is received on line 112 from the other (non-primary) memory module (or it times out). An interrupt is posted if any error was detected in either status field (lines 33-1 or 33-2), or if the non-primary memory times out.

In case two, for a read transfer, it is assumed that a data error is indicated in the status lines 33 from the primary memory or that no response is received from the primary memory. The CPUs will wait for an acknowledge from the other memory, and if no data errors are found in status bits from the other memory, circuitry of the bus interface 56 forces a change in ownership (primary memory status), then a retry is instituted to see if data is correctly read from the new primary. If good status is received from the new primary, then the stall is ended as before, and an interrupt is posted to update the system (to note one memory bad and different memory is primary). However, if data error or timeout results from this attempt to read from the new primary, then an interrupt is asserted to the processor 40 via control bus 55 and 43.

For write transfers, with the write buffer 50 bypassed, case one is where no data errors are indicated in status bits 33-1 or 33-2 from the either memory module. The stall is ended to allow instruction execution to continue. Again, an interrupt is posted if any error was detected in either status field.

For write transfers, write buffer 50 bypassed, case two is where a data error is indicated in status from the primary memory, or no response is received from the primary memory. The interface controller of each CPU waits for an acknowledge from the other memory module, and if no data errors are found in the status from the other memory an ownership change is forced and an interrupt is posted. But if data errors or timeout occur for the other (new primary) memory module, then an interrupt is asserted to the processor 40.

For write transfers with the write buffer 50 enabled so the CPU chip is not stalled by a write operation, case one is with no errors indicated in status from either memory module. The transfer is ended, so another bus transfer can begin. But if

any error is detected in either status field an interrupt is posted.

For write transfers, write buffer 50 enabled, case two is where a data error is indicated in status from the primary memory, or no response is received from the primary memory. The mechanism waits for an acknowledge from the other memory, and if no data error is found in the status from the other memory then an ownership change is forced and an interrupt is posted. But if data error or timeout occur for the other memory, then an interrupt is posted.

Once it has been determined by the mechanism just described that a memory module 14 or 15 is faulty, the fault condition is signalled to the operator, but the system can continue operating. The operator will probably wish to replace the memory board containing the faulty module, which can be done while the system is powered up and operating. The system is then able to re-integrate the new memory board without a shutdown. This mechanism also works to revive a memory module that failed to execute a write due to a soft error but then tested good so it need not be physically replaced. The task is to get the memory module back to a state where its data is identical to the other memory module. This revive mode is a two step process. First, it is assumed that the memory is uninitialized and may contain parity errors, so good data with good parity must be written into all locations, this could be all zeros at this point, but since all writes are executed on both memories the way this first step is accomplished is to read a location in the good memory module then write this data to the same location in both memory modules 14 and 15. This is done while ordinary operations are going on, interleaved with the task being performed. Writes originating from the I/O busses 24 or 25 are ignored by this revive routine in its first stage. After all locations have been thus written, the next step is the same as the first except that I/O accesses are also written; that is, I/O writes from the I/O busses 24 or 25 are executed as they occur in ordinary traffic in the executing task, interleaved with reading every location in the good memory and writing this same data to the same location in both memory modules. When the modules have been addressed from zero to maximum address in this second step, the memories are identical. During this second revive step, both CPUs and I/O processors expect the memory module being revived to perform all operations without errors. The I/O processors 26, 27 will not use data presented by the memory module being revived during data read transfers. After completing the revive process the revived memory can then be (if necessary) designated primary.

A similar revive process is provided for CPU

modules. When one CPU is detected faulty (as by the memory voter 100, etc.) the other two continue to operate, and the bad CPU board can be replaced without system shutdown. When the new CPU board has run its power-on self-test routines from on-board ROM 63, it signals this to the other CPUs, and a revive routine is executed. First, the two good CPUs will copy their state to global memory, then all three CPUs will execute a "soft reset" whereby the CPUs reset and start executing from their initialization routines in ROM, so they will all come up at the exact same point in their instruction stream and will be synchronized, then the saved state is copied back into all three CPUs and the task previously executing is continued.

As noted above, the vote circuit 100 in each memory module determines whether or not all three CPUs make identical memory references. If so, the memory operation is allowed to proceed to completion. If not, a CPU fault mode is entered. The CPU which transmits a different memory reference, as detected at the vote circuit 100, is identified in the status returned on bus 33-1 and or 33-2. An interrupt is posted and a software subsequently puts the faulty CPU offline. This offline status is reflected on status bus 32. The memory reference where the fault was detected is allowed to complete based upon the two-out-of-three vote, then until the bad CPU board has been replaced the vote circuit 100 requires two identical memory requests from the two good CPUs before allowing a memory reference to proceed. The system is ordinarily configured to continue operating with one CPU off-line, but not two. However, if it were desired to operate with only one good CPU, this is an alternative available. A CPU is voted faulty by the voter circuit 100 if different data is detected in its memory request, and also by a time-out; if two CPUs send identical memory requests, but the third does not send any signals for a preselected time-out period, that CPU is assumed to be faulty and is placed off-line as before.

The I/O arrangement of the system has a mechanism for software reintegration in the event of a failure. That is, the CPU and memory module core is hardware fault-protected as just described, but the I/O portion of the system is software fault-protected. When one of the I/O processors 26 or 27 fails, the controllers 30 bound to that I/O processor by software as mentioned above are switched over to the other I/O processor by software; the operating system rewrites the addresses in the I/O page table to use the new addresses for the same controllers, and from then on these controllers are bound to the other one of the pair of I/O processors 26 or 27. The error or fault can be detected by a bus error terminating a bus cycle at the bus interface 56, producing an exception dispatching into

the kernel through an exception handler routine that will determine the cause of the exception, and then (by rewriting addresses in the I/O table) move all the controllers 30 from the failed I/O processor 26 or 27 to the other one.

When the bus interface 56 detects a bus error as just described, the fault must be isolated before the reintegration scheme is used. When a CPU does a write, either to one of the I/O processors 26 or 27 or to one of the I/O controllers 30 on one of the busses 28 (e.g., to one of the control or status registers, or data registers, in one of the I/O elements), this is a bypass operation in the memory modules and both memory modules execute the operation, passing it on to the two I/O busses 24 and 25; the two I/O processors 26 and 27 both monitor the busses 24 and 25 and check parity and check the commands for proper syntax via the controllers 126. For example, if the CPUs are executing a write to a register in an I/O processor 26 or 27, if either one of the memory modules presents a valid address, valid command and valid data (as evidenced by no parity errors and proper protocol), the addressed I/O processor will write the data to the addressed location and respond to the memory module with an Acknowledge indication that the write was completed successfully. Both memory modules 14 and 15 are monitoring the responses from the I/O processor 26 or 27 (i.e., the address and data acknowledge signals of Figure 9, and associated status), and both memory modules respond to the CPUs with operation status on lines 33-1 and 33-2. (If this had been a read, only the primary memory module would return data, but both would return status.) Now the CPUs can determine if both executed the write correctly, or only one, or none. If only one returns good status, and that was the primary, then there is no need to force an ownership change, but if the backup returned good and the primary bad, then an ownership change is forced to make the one that executed correctly now the primary. In either case an interrupt is entered to report the fault. At this point the CPUs do not know whether it is a memory module or something downstream of the memory modules that is bad. So, a similar write is attempted to the other I/O processor, but if this succeeds it does not necessarily prove the memory module is bad because the I/O processor initially addressed could be hanging up a line on the bus 24 or 25, for example, and causing parity errors. So, the process can then selectively shut off the I/O processors and retry the operations, to see if both memory modules can correctly execute a write to the same I/O processor. If so, the system can continue operating with the bad I/O processor off-line until replaced and reintegrated. But if the retry still gives bad status from one memory, the memory can be off-

line, or further fault-isolation steps taken to make sure the fault is in the memory and not in some other element; this can include switching all the controllers 30 to one I/O processor 26 or 27 then issuing a reset command to the off I/O processor and retry communication with the online I/O processor with both memory modules live - then if the reset I/O processor had been corrupting the bus 24 or 25 its bus drivers will have been turned off by the reset so if the retry of communication to the online I/O processor (via both busses 24 and 25) now returns good status it is known that the reset I/O processor was at fault. In any event, for each bus error, some type of fault isolation sequence is implemented to determine which system component needs to be forced offline.

Synchronization:

The processors 40 used in the illustrative embodiment are of pipelined architecture with overlapped instruction execution, as discussed above with reference to Figures 4 and 5. Since a synchronization technique used in this embodiment relies upon cycle counting, i.e., incrementing a counter 71 and a counter 73 of Figure 2 every time an instruction is executed, generally as set forth in application Ser. No. 118,503, there must be a definition of what constitutes the execution of an instruction in the processor 40. A straightforward definition is that every time the pipeline advances an instruction is executed. One of the control lines in the control bus 43 is a signal RUN# which indicates that the pipeline is stalled; when RUN# is high the pipeline is stalled, when RUN# is low (logic zero) the pipeline advances each machine cycle. This RUN# signal is used in the numeric processor 46 to monitor the pipeline of the processor 40 so this coprocessor 46 can run in lockstep with its associated processor 40. This RUN# signal in the control bus 43 along with the clock 17 are used by the counters 71 and 73 to count Run cycles.

The size of the counter register 71, in a preferred embodiment, is chosen to be 4096, i.e., 2^{12} , which is selected because the tolerances of the crystal oscillators used in the clocks 17 are such that the drift in about 4K Run cycles on average results in a skew or difference in number of cycles run by a processor chip 40 of about all that can be reasonably allowed for proper operation of the interrupt synchronization as explained below. One synchronization mechanism is to force action to cause the CPUs to synchronize whenever the counter 71 overflows. One such action is to force a cache miss in response to an overflow signal OVFL from the counter 71; this can be done by merely generating a false Miss signal (e.g., TagValid bit

not set) on control bus 43 for the next I-cache reference, thus forcing a cache miss exception routine to be entered and the resultant memory reference will produce synchronization just as any memory reference does. Another method of forcing synchronization upon overflow of counter 71 is by forcing a stall in the processor 40, which can be done by using the overflow signal OVFL to generate a CP Busy (coprocessor busy) signal on control bus 43 via logic circuit 71a of Figure 2; this CP Busy signal always results in the processor 40 entering stall until CP Busy is deasserted. All three processors will enter this stall because they are executing the same code and will count the same cycles in their counter 71, but the actual time they enter the stall will vary; the logic circuit 71a receives the RUN# signal from bus 43 of the other two processors via input R#, so when all three have stalled the CP Busy signal is released and the processors will come out of stall in synch again.

Thus, two synchronization techniques have been described, the first being the synchronization resulting from voting the memory references in circuits 100 in the memory modules, and the second by the overflow of counter 71 as just set forth. In addition, interrupts are synchronized, as will be described below. It is important to note, however, that the processors 40 are basically running free at their own clock speed, and are substantially decoupled from one another, except when synchronizing events occur. The fact that microprocessors are used as illustrated in Figures 4 and 5 would make lock-step synchronization with a single clock more difficult, and would degrade performance; also, use of the write buffer 50 serves to decouple the processors, and would be much less effective with close coupling of the processors. Likewise, the high-performance resulting from using instruction and data caches, and virtual memory management with the TLBs 83, would be more difficult to implement if close coupling were used, and performance would suffer.

The interrupt synchronization technique must distinguish between real time and so-called "virtual time". Real time is the external actual time, clock-on-the-wall time, measured in seconds, or for convenience, measured in machine cycles which are 60-nsec divisions in the example. The clock generators 17 each produce clock pulses in real time, of course. Virtual time is the internal cycle-count time of each of the processor chips 40 as measured in each one of the cycle counters 71 and 73, i.e., the instruction number of the instruction being executed by the processor chip, measured in instructions since some arbitrary beginning point. Referring to Figure 10, the relationship between real time, shown as t_0 to t_{12} , and virtual time, shown as instruction number (modulo-16 count in count reg-

ister 73) I_0 to I_{15} , is illustrated. Each row of Figure 10 is the cycle count for one of the CPUs A, B or C, and each column is a "point" in real time. The clocks for the CPUs will most likely be out of phase, so the actual time correlation will be as seen in Figure 10a, where the instruction numbers (columns) are not perfectly aligned, i.e., the cycle-count does not change on aligned real-time machine cycle boundaries; however, for explanatory purposes the illustration of Figure 10 will suffice. In Figure 10, at real time t_3 the CPU-A is at the third instruction, CPU-B is at count-9 or executing the ninth instruction, and CPU-C is at the fourth instruction. Note that both real time and virtual time can only advance.

The processor chip 40 in a CPU stalls under certain conditions when a resource is not available, such as a D-cache 45 or I-cache 44 miss during a load or an instruction fetch, or a signal that the write buffer 50 is full during a store operation, or a "CP Busy" signal via the control bus 43 that the coprocessor 46 is busy (the coprocessor receives an instruction it cannot yet handle due to data dependency or limited processing resources), or the multiplier/divider 79 is busy (the internal multiply/divide circuit has not completed an operation at the time the processor attempts to access the result register). Of these, the caches 44 and 45 are "passive resources" which do not change state without intervention by the processor 40, but the remainder of the items are active resources that can change state while the processor is not doing anything to act upon the resource. For example, the write buffer 50 can change from full to empty with no action by the processor (so long as the processor does not perform another store operation). So there are two types of stalls: stalls on passive resources and stalls on active resources. Stalls on active resources are called interlock stalls.

Since the code streams executing on the CPUs A, B and C are the same, the states of the passive resources such as caches 44 and 45 in the three CPUs are necessarily the same at every point in virtual time. If a stall is a result of a conflict at a passive resource (e.g., the data cache 45) then all three processors will perform a stall, and the only variable will be the length of the stall. Referring to Figure 11, assume the cache miss occurs at I_4 , and that the access to the global memory 14 or 15 resulting from the miss takes eight clocks (actually it may be more than eight). In this case, CPU-C begins the access to global memory 14 and 15 at t_1 , and the controller 117 for global memory begins the memory access when the first processor CPU-C signals the beginning of the memory access. The controller 117 completes the access eight clocks later, at t_8 , although CPU-A and CPU-B each stalled less than the eight clocks required for

the memory access. The result is that the CPUs become synchronized in real time as well as in virtual time. This example also illustrates the advantage of overlapping the access to DRAM 104 and the voting in circuit 100.

Interlock stalls present a different situation from passive resource stalls. One CPU can perform an interlock stall when another CPU does not stall at all. Referring to Figure 12, an interlock stall caused by the write buffer 50 is illustrated. The cycle-counts for CPU-A and CPU-B are shown, and the full flags A_{wb} and B_{wb} from write buffers 50 for CPU-A and CPU-B are shown below the cycle-counts (high or logic one means full, low or logic zero means empty). The CPU checks the state of the full flag every time a store operation is executed; if the full flag is set, the CPU stalls until the full flag is cleared then completes the store operation. The write buffer 50 sets the full flag if the store operation fills the buffer, and clears the full flag whenever a store operation drains one word from the buffer thereby freeing a location for the next CPU store operation. At time t_0 the CPU-B is three clocks ahead of CPU-A, and the write buffers are both full. Assume the write buffers are performing a write operation to global memory, so when this write completes during t_5 the write buffer full flags will be cleared; this clearing will occur synchronously in t_6 in real time (for the reason illustrated by Figure 11) but not synchronously in virtual time. Now, assume the instruction at cycle-count I_6 is a store operation; CPU-A executes this store at t_6 after the write buffer full flag is cleared, but CPU-B tries to execute this store operation at t_3 and finds the write buffer full flag is still set and so has to stall for three clocks. Thus, CPU-B performs a stall that CPU-A did not.

The property that one CPU may stall and the other not stall imposes a restriction on the interpretation of the cycle counter 71. In Figure 12, assume interrupts are presented to the CPUs on a cycle count of I_7 (while the CPU-B is stalling from the I_6 instruction). The run cycle for cycle count I_7 occurs for both CPUs at t_7 . If the cycle counter alone presents the interrupt to the CPU, then CPU-A would see the interrupt on cycle count I_7 but CPU-B would see the interrupt during a stall cycle resulting from cycle count I_6 , so this method of presenting interrupts would cause the two CPUs to take an exception on different instructions, a condition that would not have occurred if either all of the CPUs stalled or none stalled.

Another restriction on the interpretation of the cycle counter is that there should not be any delays between detecting the cycle count and performing an action. Again referring to Figure 12, assume interrupts are presented to the CPUs on cycle count I_6 , but because of implementation re-

strictions an extra clock delay is interposed between detection of cycle count I_6 and presentation of the interrupt to the CPU. The result is that CPU-A sees this interrupt on cycle count I_7 , but CPU-B will see the interrupt during the stall from cycle count I_6 , causing the two CPUs to take an exception on different instructions. Again, the importance of monitoring the state of the instruction pipeline in real time is illustrated.

Interrupt Synchronization:

The three CPUs of the system of Figures 1-3 are required to function as a single logical processor, thus requiring that the CPUs adhere to certain restrictions regarding their internal state to ensure that the programming model of the three CPUs is that of a single logical processor. Except in failure modes and in diagnostic functions, the instruction streams of the three CPUs are required to be identical. If not identical, then voting global memory accesses at voting circuitry 100 of Figure 6 would be difficult; the voter would not know whether one CPU was faulty or whether it was executing a different sequence of instructions. The synchronization scheme is designed so that if the code stream of any CPU diverges from the code stream of the other CPUs, then a failure is assumed to have occurred. Interrupt synchronization provides one of the mechanisms of maintaining a single CPU image.

All interrupts are required to occur synchronous to virtual time, ensuring that the instruction streams of the three processors CPU-A, CPU-B and CPU-C will not diverge as a result of interrupts (there are other causes of divergent instruction streams, such as one processor reading different data than the data read by the other processors). Several scenarios exist whereby interrupts occurring asynchronous to virtual time would cause the code streams to diverge. For example, an interrupt causing a context switch on one CPU before process A completes, but causing the context switch after process A completes on another CPU would result in a situation where, at some point later, one CPU continues executing process A, but the other CPU cannot execute process A because that process had already completed. If in this case the interrupts occurred asynchronous to virtual time, then just the fact that the exception program counters were different could cause problems. The act of writing the exception program counters to global memory would result in the voter detecting different data from the three CPUs, producing a vote fault.

Certain types of exceptions in the CPUs are inherently synchronous to virtual time. One example is a breakpoint exception caused by the execu-

tion of a breakpoint instruction. Since the instruction streams of the CPUs are identical, the breakpoint exception occurs at the same point in virtual time on all three of the CPUs. Similarly, all such internal exceptions inherently occur synchronous to virtual time. For example, TLB exceptions are internal exceptions that are inherently synchronous. TLB exceptions occur because the virtual page number does not match any of the entries in the TLB 83. Because the act of translating addresses is solely a function of the instruction stream (exactly as in the case of the breakpoint exception), the translation is inherently synchronous to virtual time. In order to ensure that TLB exceptions are synchronous to virtual time, the state of the TLBs 83 must be identical in all three of the CPUs 11, 12 and 13, and this is guaranteed because the TLB 83 can only be modified by software. Again, since all of the CPUs execute the same instruction stream, the state of the TLBs 83 are always changed synchronous to virtual time. So, as a general rule of thumb, if an action is performed by software then the action is synchronous to virtual time. If an action is performed by hardware, which does not use the cycle counters 71, then the action is generally synchronous to real time.

External exceptions are not inherently synchronous to virtual time. I/O devices 26, 27 or 30 have no information about the virtual time of the three CPUs 11, 12 and 13. Therefore, all interrupts that are generated by these I/O devices must be synchronized to virtual time before presenting to the CPUs, as explained below. Floating point exceptions are different from I/O device interrupts because the floating point coprocessor 46 is tightly coupled to the microprocessor 40 within the CPU.

External devices view the three CPUs as one logical processor, and have no information about the synchronicity or lack of synchronicity between the CPUs, so the external devices cannot produce interrupts that are synchronous with the individual instruction stream (virtual time) of each CPU. Without any sort of synchronization, if some external device drove an interrupt at time real time t_1 of Figure 10, and the interrupt was presented directly to the CPUs at this time then the three CPUs would take an exception trap at different instructions, resulting in an unacceptable state of the three CPUs. This is an example of an event (assertion of an interrupt) which is synchronous to real time but not synchronous to virtual time.

Interrupts are synchronized to virtual time in the system of Figures 1-3 by performing a distributed vote on the interrupts and then presenting the interrupt to the processor on a predetermined cycle count. Figure 13 shows a more detailed block diagram of the interrupt synchronization logic 65 of Figure 2. Each CPU contains a distributor 135

which captures the external interrupt from the line 69 or 70 coming from the modules 14 or 15; this capture occurs on a predetermined cycle count, e.g., at count-4 as signalled on an input line CC-4 from the counter 71. The captured interrupt is distributed to the other two CPUs via the inter-CPU bus 18. These distributed interrupts are called pending interrupts. There are three pending interrupts, one from each CPU 11, 12 and 13. A voter circuit 136 captures the pending interrupts and performs a vote to verify that all of the CPUs did receive the external interrupt request. On a predetermined cycle count (detected from the cycle counter 71), in this example cycle-8 received by input line CC-8, the interrupt voter 136 presents the interrupt to the interrupt pin on its respective microprocessor 40 via line 137 and control bus 55 and 43. Since the cycle count that is used to present the interrupt is predetermined, all of the microprocessors 40 will receive the interrupt on the same cycle count and thus the interrupt will have been synchronized to virtual time.

Figure 14 shows the sequence of events for synchronizing interrupts to virtual time. The rows labeled CPU-A, CPU-B, and CPU-C indicate the cycle count in counter 71 of each CPU at a point in real time. The rows labeled IRQ_A_PEND, IRQ_B_PEND, and IRQ_C_PEND indicate the state of the interrupt pending bits coupled via the inter-CPU bus 18 to the input of the voters 136 (a one signifies that the pending bit is set). The rows labeled IRQ_A, IRQ_B, and IRQ_C indicate the state of the interrupt input pin on the microprocessor 40 (the signals on lines 137), where a one signifies that an interrupt is present at the input pin. In Figure 14, the external interrupt (EX_IRQ) is asserted on line 69 at t_0 . If the interrupt distributor 135 captures and then distributes the interrupt to the inter-CPU bus 18 on cycle count 4, then IRQ_C_PEND will go active at t_1 , IRQ_B_PEND will go active at t_2 , and IRQ_A_PEND will go active at t_4 . If the interrupt voter 136 captures and then votes the interrupt pending bits on cycle count 8, then IRQ_C will go active at t_5 , IRQ_B will go active at t_5 , and IRQ_A will go active at t_8 . The result is that the interrupts were presented to the CPUs at different points in real time but at the same point in virtual time (i.e. cycle count 8).

Figure 15 illustrates a scenario which requires the algorithm presented in Figure 14 to be modified. Note that the cycle counter 71 is here represented by a modulo 8 counter. The external interrupt (EX_IRQ) is asserted at time t_3 , and the interrupt distributor 135 captures and then distributes the interrupt to the inter-CPU bus 18 on cycle count 4. Since CPU-B and CPU-C have executed cycle count 4 before time t_3 , their interrupt distributor does not capture the external interrupt. CPU-A,

however, executes cycle count 4 after time t_3 . The result is that CPU-A captures and distributes the external interrupt at time t_4 . But if the interrupt voter 136 captures and votes the interrupt pending bits on cycle 7, the interrupt voter on CPU-A captures the IRQ_A_PEND signal at time t_7 , when the two other interrupt pending bits are not set. The interrupt voter 136 on CPU-A recognizes that not all of the CPUs have distributed the external interrupt and thus places the captured interrupt pending bit in a holding register 138. The interrupt voters 136 on CPU-B and CPU-C capture the single interrupt pending bit at times t_5 and t_4 respectively. Like the interrupt voter on CPU-A, the voters recognize that not all of the interrupt pending bits are set, and thus the single interrupt pending bit that is set is placed into the holding register 138. When the cycle counter 71 on each CPU reaches a cycle count of 7, the counter rolls over and begins counting at cycle count 0. Since the external interrupt is still asserted, the interrupt distributor 135 on CPU-B and CPU-C will capture the external interrupt at times t_{10} and t_9 respectively. These times correspond to when the cycle count becomes equal to 4. At time t_{12} , the interrupt voter on CPU-C captures the interrupt pending bits on the inter-CPU bus 18. The voter 136 determines that all of the CPUs did capture and distribute the external interrupt and thus presents the interrupt to the processor chip 40. At times t_{13} and t_{15} , the interrupt voters 136 on CPU-B and CPU-A capture the interrupt pending bits and then presents the interrupt to the processor chip 40. The result is that all of the processor chips received the external interrupt request at identical instructions, and the information saved in the holding registers is not needed.

Holding Register:

In the interrupt scenario presented above with reference to Figure 15, the voter 136 uses a holding register 138 to save some state information. In particular, the saved state was that some, but not all, of the CPUs captured and distributed an external interrupt. If the system does not have any faults (as was the situation in Figure 15) then this state information is not necessary because, as shown in the previous example, external interrupts can be synchronized to virtual time without the use of the holding register 138. The algorithm is that the interrupt voter 136 captures and votes the interrupt pending bits on a predetermined cycle count. When all of the interrupt pending bits are asserted, then the interrupt is presented to the processor chip 40 on the predetermined cycle count. In the example of Figure 15, the interrupts were voted on cycle count 7.

Referring to Figure 15, if CPU-C fails and the

failure mode is such that the interrupt distributor 135 does not function correctly, then if the interrupt voters 136 waited until all of the interrupt pending bits were set before presenting the interrupt to the processor chip 40, the result would be that the interrupt would never get presented. Thus, a single fault on a single CPU renders the entire interrupt chain on all of the CPUs inoperable.

The holding register 138 provides a mechanism for the voter 136 to know that the last interrupt vote cycle captured at least one, but not all, of the interrupt pending bits. The interrupt vote cycle occurs on the cycle count that the interrupt voter captures and votes the interrupt pending bits. There are only two scenarios that result in some of the interrupt pending bits being set. One is the scenario presented in reference to Figure 15 in which the external interrupt is asserted before the interrupt distribution cycle on some of the CPUs but after the interrupt distribution cycle on other CPUs. In the second scenario, at least one of the CPUs fails in a manner that disables the interrupt distributor. If the reason that only some of the interrupt pending bits are set at the interrupt vote cycle is case one scenario, then the interrupt voter is guaranteed that all of the interrupt pending bits will be set on the next interrupt vote cycle. Therefore, if the interrupt voter discovers that the holding register has been set and not all of the interrupt pending bits are set, then an error must exist on one or more of the CPUs. This assumes that the holding register 138 of each CPU gets cleared when an interrupt is serviced, so that the state of the holding register does not represent stale state on the interrupt pending bits. In the case of an error, the interrupt voter 136 can present the interrupt to the processor chip 40 and simultaneously indicate that an error has been detected in the interrupt synchronization logic.

The interrupt voter 136 does not actually do any voting but instead merely checks the state of the interrupt pending bits and the holding register 137 to determine whether or not to present an interrupt to the processor chip 40 and whether or not to indicate an error in the interrupt logic.

Modulo Cycle Counters:

The interrupt synchronization example of Figure 15 represented the interrupt cycle counter 71 as a modulo N counter (e.g., a modulo 8 counter). Using a modulo N cycle counter simplified the description of the interrupt voting algorithm by allowing the concept of an interrupt vote cycle. With a modulo N cycle counter, the interrupt vote cycle can be described as a single cycle count which lies between 0 and N-1 where N is the modulo of the cycle counter. Whatever value of cycle counter is

chosen for the interrupt vote cycle, that cycle count is guaranteed to occur every N cycle counts; as illustrated in Figure 15 for a modulo 8 counter, every eight counts an interrupt vote cycle occurs. The interrupt vote cycle is used here merely to illustrate the periodic nature of a modulo N cycle counter. Any event that is keyed to a particular cycle count of a modulo N cycle counter is guaranteed to occur every N cycle counts. Obviously, an infinite (i.e., non-repeating counter 71) couldn't be used.

A value of N is chosen to maximize system parameters that have a positive effect on the system and to minimize system parameters that have a negative effect on the system. Some of such effects are developed empirically. First, some of the parameters will be described; C_v and C_d are the interrupt vote cycle and the interrupt distribution cycle respectively (in the circuit of Figure 13 these are the inputs CC-8 and CC-4, respectively). The value of C_v and C_d must lie in the range between 0 and N-1 where N is the modulo of the cycle counter. D_{max} is the maximum amount of cycle count drift between the three processors CPU-A, -B and -C that can be tolerated by the synchronization logic. The processor drift is determined by taking a snapshot of the cycle counter 71 from each CPU at a point in real time. The drift is calculated by subtracting the cycle count of the slowest CPU from the cycle count of the fastest CPU, performed as modulo N subtraction. The value of D_{max} is described as a function of N and the values of C_v and C_d .

First, D_{max} will be defined as a function of the difference $C_v - C_d$, where the subtraction operation is performed as modulo N subtraction. This allows us to choose values of C_v and C_d that maximize D_{max} . Consider the scenario in Figure 16. Suppose that $C_d = 8$ and $C_v = 9$. From Figure 16 the processor drift can be calculated to be $D_{max} = 4$. The external interrupt on line 69 is asserted at time t_4 . In this case, CPU-B will capture and distribute the interrupt at time t_5 . CPU-B will then capture and vote the interrupt pending bits at time t_6 . This scenario is inconsistent with the interrupt synchronization algorithm presented earlier because CPU-B executes its interrupt vote cycle before CPU-A has performed the interrupt distribution cycle. The flaw with this scenario is that the processors have drifted further apart than the difference between C_v and C_d . The relationship can be formally written as

$$\text{Equation (1)} \quad C_v - C_d < D_{max} - e$$

where e is the time needed for the interrupt pending bits to propagate on the inter-CPU bus 18. In previous examples, e has been assumed to be zero. Since wall-clock time has been quantized in

clock cycle (Run cycle) increments, e can also be quantized. Thus the equation becomes

$$\text{Equation (2)} \quad C_v - C_d < D_{\max} - 1$$

where D_{\max} is expressed as an integer number of cycle counts.

Next, the maximum drift can be described as a function of N . Figure 17 illustrates a scenario in which $N = 4$ and the processor drift $D = 3$. Suppose that $C_d = 0$. The subscripts on cycle count 0 of each processor denote the quotient part (Q) of the instruction cycle count. Since the cycle count is now represented in modulo N , the value of the cycle counter is the remainder portion of I/N where I is the number of instructions that have been executed since time t_0 . The Q of the instruction cycle count is the integer portion of I/N . If the external interrupt is asserted at time t_3 , then CPU-A will capture and distribute the interrupt at time t_4 , and CPU-B will execute its interrupt distribution cycle at time t_5 . This presents a problem because the interrupt distribution cycle for CPU-A has $Q = 1$ and the interrupt distribution cycle for CPU-B has $Q=2$. The synchronization logic will continue as if there are no problems and will thus present the interrupt to the processors on equal cycle counts. But the interrupt will be presented to the processors on different instructions because the Q of each processor is different. The relationship of D_{\max} as a function of N is therefore

$$\text{Equation (3)} \quad N/2 > D_{\max}$$

where N is an even number and D_{\max} is expressed as an integer number of cycle counts. (These equations 2 and 3 can be shown to be both equivalent to the Nyquist theorem in sampling theory.) Combining equations 2 and 3 gives

$$\text{Equation (4)} \quad C_v - C_d < N/2 - 1$$

which allows optimum values of C_v and C_d to be chosen for a given value of N . All of the above equations suggest that N should be as large as possible. The only factor that tries to drive N to a small number is interrupt latency. Interrupt latency is the time interval between the assertion of the external interrupt on line 69 and the presentation of the interrupt to the microprocessor chip on line 137. Which processor should be used to determine the interrupt latency is not a clear-cut choice. The three microprocessors will operate at different speeds because of the slight differences in the crystal oscillators in clock sources 17 and other factors. There will be a fastest processor, a slowest processor, and the other processor. Defining the interrupt latency with respect to the slowest pro-

cessor is reasonable because the performance of system is ultimately determined by the performance of the slowest processor. The maximum interrupt latency is

$$\text{Equation (5)} \quad L_{\max} = 2N - 1$$

where L_{\max} is the maximum interrupt latency expressed in cycle counts. The maximum interrupt latency occurs when the external interrupt is asserted after the interrupt distribution cycle C_d of the fastest processor but before the interrupt distribution cycle C_d of the slowest processor. The calculation of the average interrupt latency L_{ave} is more complicated because it depends on the probability that the external interrupt occurs after the interrupt distribution cycle of the fastest processor and before the interrupt distribution cycle of the slowest processor. This probability depends on the drift between the processors which in turn is determined by a number of external factors. If we assume that these probabilities are zero, then the average latency may be expressed as

$$\text{Equation (6)} \quad L_{\text{ave}} = N/2 + (C_v - C_d)$$

Using these relationships, values of N , C_v , and C_d are chosen using the system requirements for D_{\max} and interrupt latency. For example, choosing $N = 128$ and $(C_v - C_d) = 10$, $L_{\text{ave}} = 74$ or about 4.4 microsec (with no stall cycles). Using the preferred embodiment where a four bit (four binary stage) counter 71a is used as the interrupt synch counter, and the distribute and vote outputs are at CC-4 and CC-8 as discussed, it is seen that $N=16$, $C_v=8$ and $C_d=4$, so $L_{\text{ave}}=16/2 + (8A) = 12$ -cycles or 0.7 microsec.

Refresh Control for Local Memory:

The refresh counter 72 counts non-stall cycles (not machine cycles) just as the counters 71 and 71a count. The object is that the refresh cycles will be introduced for each CPU at the same cycle count, measured in virtual time rather than real time. Preferably, each one of the CPUs will interpose a refresh cycle at the same point in the instruction stream as the other two. The DRAMs in local memory 16 must be refreshed on a 512 cycles per 8-msec. schedule just as mentioned above regarding the DRAMs 104 of the global memory. Thus, the counter 72 could issue a refresh command to the DRAMs 16 once every 15-microsec., addressing one row of 512, so the refresh specification would be satisfied; if a memory operation was requested during refresh then a Busy response would result until refresh was finished. But letting each CPU handle its own local

memory refresh in real time independently of the others could cause the CPUs to get out of synch, and so additional control is needed. For example, if refresh mode is entered just as a divide operation is beginning, then timing is such that one CPU could take two clocks longer than others. Or, if a non-interruptable sequence was entered by a faster CPU then the others went into refresh before entering this routine, the CPUs could walk away from one another. However, using the cycle counter 71 (instead of real time) to avoid some of these problems means that stall cycles are not counted, and if a loop is entered causing many stalls (some can cause a 7-to-1 stall-to-run ratio) then the refresh specification is not met unless the period is decreased substantially from the 15-microsec figure, but that would degrade performance. For this reason, stall cycles are also counted in a second counter 72a, seen in Figure 2, and every time this counter reaches the same number as that counted in the refresh counter 72, an additional refresh cycle is introduced. For example, the refresh counter 72 counts 2^8 or 256 Run cycles, in step with the counter 71, and when it overflows a refresh is signalled via control bus 43. Meanwhile, counter 72a counts 2^8 stall cycles (responsive to the RUN# signal and clock 17), and every time it overflows a second counter 72b is incremented (counter 72b may be merely bits 9-to-11 for the eight-bit counter 72a), so when a refresh mode is finally entered the CPU does a number of additional refreshes indicated by the number in the counter register 72b. Thus, if a long period of stall-intensive execution is encountered, the average number of refreshes will stay in the one per 15-microsec range, even if up to 7×256 stall cycles are interposed, because when finally going into a refresh mode the number of rows refreshed will catch up to the nominal refresh rate, yet there is no degradation of performance by arbitrarily shortening the refresh cycle.

Memory Management:

The CPUs 11, 12 and 13 of Figures 1-3 have memory space organized as illustrated in Figure 18. Using the example that the local memory 16 is 8-MByte and the global memory 14 or 15 is 32-MByte, note that the local memory 16 is part of the same continuous zero-to-40M map of CPU memory access space, rather than being a cache or a separate memory space; realizing that the 0-8M section is triplicated (in the three CPU modules), and the 8-40M section is duplicated, nevertheless logically there is merely a single 0-40M physical address space. An address over 8-MByte on bus 54 causes the bus interface 56 to make a request to the memory modules 14 and 15, but an address under 8-MByte will access the local memory 16

within the CPU module itself. Performance is improved by placing more of the memory used by the applications being executed in local memory 16, and so as memory chips are available in higher densities at lower cost and higher speeds, additional local memory will be added, as well as additional global memory. For example, the local memory might be 32-MByte and the global memory 128-MByte. On the other hand, if a very minimum-cost system is needed, and performance is not a major determining factor, the system can be operated with no local memory, all main memory being in the global memory area (in memory modules 14 and 15), although the performance penalty is high for such a configuration.

The content of local memory portion 141 of the map of Figure 18 is identical in the three CPUs 11, 12 and 13. Likewise, the two memory modules 14 and 15 contain identically the same data in their space 142 at any given instant. Within the local memory portion 141 is stored the kernel 143 (code) for the Unix operating system, and this area is physically mapped within a fixed portion of the local memory 16 of each CPU. Likewise, kernel data is assigned a fixed area 144 in each local memory 16; except upon boot-up, these blocks do not get swapped to or from global memory or disk. Another portion 145 of local memory 16 is employed for user program (and data) pages, which are swapped to area 146 of the global memory 14 and 15 under control of the operating system. The global memory area 142 is used as a staging area for user pages in area 146, and also as a disk buffer in an area 147; if the CPUs are executing code which performs a write of a block of data or code from local memory 16 to disk 148, then the sequence is to always write to a disk buffer area 147 instead because the time to copy to area 147 is negligible compared to the time to copy directly to the I/O processor 26 and 27 and thus via I/O controller 30 to disk 148. Then, while the CPUs proceed to execute other code, the write-to-disk operation is done, transparent to the CPUs, to move the block from area 147 to disk 148. In a like manner, the global memory area 146 is mapped to include an I/O staging 149 area, for similar treatment of I/O accesses other than disk (e.g., video).

The physical memory map of Figure 18 is correlated with the virtual memory management system of the processor 40 in each CPU. Figure 19 illustrates the virtual address map of the R2000 processor chip used in the example embodiment, although it is understood that other microprocessor chips supporting virtual memory management with paging and a protection mechanism would provide corresponding features.

In Figure 19, two separate 2-GByte virtual address spaces 150 and 151 are illustrated; the pro-

processor 40 operates in one of two modes, user mode and kernel mode. The processor can only access the area 150 in the user mode, or can access both the areas 150 and 151 in the kernel mode. The kernel mode is analogous to the supervisory mode provided in many machines. The processor 40 is configured to operate normally in the user mode until an exception is detected forcing it into the kernel mode, where it remains until a restore from exception (RFE) instruction is executed. The manner in which the memory addresses are translated or mapped depends upon the operating mode of the microprocessor, which is defined by a bit in a status register. When in the user mode, a single, uniform virtual address space 150 referred to as "kuseg" of 2-GByte size is available. Each virtual address is also extended with a 6-bit process identifier (PID) field to form unique virtual addresses for up to sixty-four user processes. All references to this segment 150 in user mode are mapped through the TLB 83, and use of the caches 144 and 145 is determined by bit settings for each page entry in the TLB entries; i.e., some pages may be cachable and some not as specified by the programmer.

When in the kernel mode, the virtual address space includes both the areas 150 and 151 of Figure 19, and this space has four separate segments kuseg 150, kseg0 152, kseg1 153 and kseg2 154. The kuseg 150 segment for the kernel mode is 2-GByte in size, coincident with the "kuseg" of the user mode, so when in the kernel mode the processor treats references to this segment just like user mode references, thus streamlining kernel access to user data. The kuseg 150 is used to hold user code and data, but the operating system often needs to reference this same code or data. The kseg0 area 152 is a 512-MByte kernel physical address space direct-mapped onto the first 512-MBytes of physical address space, and is cached but does not use the TLB 83; this segment is used for kernel executable code and some kernel data, and is represented by the area 143 of Figure 18 in local memory 16. The kseg1 area 153 is also directly mapped into the first 512-MByte of physical address space, the same as kseg0, and is uncached and uses no TLB entries. Kseg1 differs from kseg0 only in that it is uncached. Kseg1 is used by the operating system for I/O registers, ROM code and disk buffers, and so corresponds to areas 147 and 149 of the physical map of Figure 18. The kseg2 area 154 is a 1-GByte space which, like kuseg, uses TLB 83 entries to map virtual addresses to arbitrary physical ones, with or without caching. This kseg2 area differs from the kuseg area 150 only in that it is not accessible in the user mode, but instead only in the kernel mode. The operating system uses kseg2 for stacks and per-

process data that must remap on context switches, for user page tables (memory map), and for some dynamically-allocated data areas. Kseg2 allows selective caching and mapping on a per page basis, rather than requiring an all-or-nothing approach.

The 32-bit virtual addresses generated in the registers 76 or PC 80 of the microprocessor chip and output on the bus 84 are represented in Figure 20, where it is seen that bits 0-11 are the offset used unconditionally as the low-order 12-bits of the address on bus 42 of Figure 3, while bits 12-31 are the VPN or virtual page number in which bits 29-31 select between kuseg, kseg0, kseg1 and kseg2. The process identifier PID for the currently-executing process is stored in a register also accessible by the TLB. The 64-bit TLB entries are represented in Figure 20 as well, where it is seen that the 20-bit VPN from the virtual address is compared to the 20-bit VPN field located in bits 44-63 of the 64-bit entry, while at the same time the PID is compared to bits 38-43; if a match is found in any of the sixty-four 64-bit TLB entries, the page frame number PFN at bits 12-31 of the matched entry is used as the output via busses 82 and 42 of Figure 3 (assuming other criteria are met). Other one-bit values in a TLB entry include N, D, V and G. N is the non-cachable indicator, and if set the page is non-cachable and the processor directly accesses local memory or global memory instead of first accessing the cache 44 or 45. D is a write-protect bit, and if set means that the location is "dirty" and therefore writable, but if zero a write operation causes a trap. The V bit means valid if set, and allows the TLB entries to be cleared by merely resetting the valid bits; this V bit is used in the page-swapping arrangement of this system to indicate whether a page is in local or global memory. The G bit is to allow global accesses which ignore the PID match requirement for a valid TLB translation; in kseg2 this allows the kernel to access all mapped data without regard for PID.

The device controllers 30 cannot do DMA into local memory 16 directly, and so the global memory is used as a staging area for DMA type block transfers, typically from disk 148 or the like. The CPUs can perform operations directly at the controllers 30, to initiate or actually control operations by the controllers (i.e., programmed I/O), but the controllers 30 cannot do DMA except to global memory; the controllers 30 can become the VMEbus (bus 28) master and through the I/O processor 26 or 27 do reads or writes directly to global memory in the memory modules 14 and 15.

Page swapping between global and local memories (and disk) is initiated either by a page fault or by an aging process. A page fault occurs when a process is executing and attempts to ex-

ecute from or access a page that is in global memory or on disk; the TLB 83 will show a miss and a trap will result, so low level trap code in the kernel will show the location of the page, and a routine will be entered to initiate a page swap. If the page needed is in global memory, a series of commands are sent to the DMA controller 74 to write the least-recently-used page from local memory to global memory and to read the needed page from global to local. If the page is on disk, commands and addresses (sectors) are written to the controller 30 from the CPU to go to disk and acquire the page, then the process which made the memory reference is suspended. When the disk controller has found the data and is ready to send it, an interrupt is signalled which will be used by the memory modules (not reaching the CPUs) to allow the disk controller to begin a DMA to global memory to write the page into global memory, and when finished the CPU is interrupted to begin a block transfer under control of DMA controller 74 to swap a least used page from local to global and read the needed page to local. Then, the original process is made runnable again, state is restored, and the original memory reference will again occur, finding the needed page in local memory. The other mechanism to initiate page swapping is an aging routine by which the operating system periodically goes through the pages in local memory marking them as to whether or not each page has been used recently, and those that have not are subject to be pushed out to global memory. A task switch does not itself initiate page swapping, but instead as the new task begins to produce page faults pages will be swapped as needed, and the candidates for swapping out are those not recently used.

If a memory reference is made and a TLB miss is shown, but the page table lookup resulting from the TLB miss exception shows the page is in local memory, then a TLB entry is made to show this page to be in local memory. That is, the process takes an exception when the TLB miss occurs, goes to the page tables (in the kernel data section), finds the table entry, writes to TLB, then the process is allowed to proceed. But if the memory reference shows a TLB miss, and the page tables show the corresponding physical address is in global memory (over 8M physical address), the TLB entry is made for this page, and when the process resumes it will find the page entry in the TLB as before; yet another exception is taken because the valid bit will be zero, indicating the page is physically not in local memory, so this time the exception will enter a routine to swap the page from global to local and validate the TLB entry, so execution can then proceed. In the third situation, if the page tables show address for the memory

reference is on disk, not in local or global memory, then the system operates as indicated above, i.e., the process is put off the run queue and put in the sleep queue, a disk request is made, and when the disk has transferred the page to global memory and signalled a command-complete interrupt, then the page is swapped from global to local, and the TLB updated, then the process can execute again.

Private Memory:

Although the memory modules 14 and 15 store the same data at the same locations, and all three CPUs 11, 12 and 13 have equal access to these memory modules, there is a small area of the memory assigned under software control as a private memory in each one of the memory modules. For example, as illustrated in Figure 21, an area 155 of the map of the memory module locations is designated the private memory area, and is writable only when the CPUs issue a "private memory write" command on bus 59. In an example embodiment, the private memory area 155 is a 4K page starting at the address contained in a register 156 in the bus interface 56 of each one of the CPU modules; this starting address can be changed under software control by writing to this register 156 by the CPU. The private memory area 155 is further divided between the three CPUs; only CPU-A can write to area 155a, CPU-B to area 155b, and CPU-C to area 155c. One of the command signals in bus 57 is set by the bus interface 56 to inform the memory modules 14 and 15 that the operation is a private write, and this is set in response to the address generated by the processor 40 from a Store instruction; bits of the address (and a Write command) are detected by a decoder 157 in the bus interface (which compares bus addresses to the contents of register 156) and used to generate the "private memory write" command for bus 57. In the memory module, when a write command is detected in the registers 94, 95 and 96, and the addresses and commands are all voted good (i.e., in agreement) by the vote circuit 100, then the control circuit 100 allows the data from only one of the CPUs to pass through to the bus 101, this one being determined by two bits of the address from the CPUs. During this private write, all three CPUs present the same address on their bus 57 but different data on their bus 58 (the different data is some state unique to the CPU, for example). The memory modules vote the addresses and commands, and select data from only one CPU based upon part of the address field seen on the address bus. To allow the CPUs to vote some data, all three CPUs will do three private writes (there will be three writes on the busses 21, 22 and 23) of some state information unique to a CPU, into both mem-

ory modules 14 and 15. During each write, each CPU sends its unique data, but only one is accepted each time. So, the software sequence executed by all three CPUs is (1) Store (to location 155a), (2) Store (to location 155b), (3) Store (to location 155c). But data from only one CPU is actually written each time, and the data is not voted (because it is or could be different and could show a fault if voted). Then, the CPUs can vote the data by having all three CPUs read all three of the locations 155a, 155b and 155c, and by software compare this data. This type of operation is used in diagnostics, for example, or in interrupts to vote the cause register data.

The private-write mechanism is used in fault detection and recovery. For example, if the CPUs detect a bus error upon making a memory read request, such as a memory module 14 or 15 returning bad status on lines 33-1 or 33-2. At this point a CPU doesn't know if the other CPUs received the same status from the memory module; the CPU could be faulty or its status detection circuit faulty, or, as indicated, the memory could be faulty. So, to isolate the fault, when the bus fault routine mentioned above is entered, all three CPUs do a private write of the status information they just received from the memory modules in the preceding read attempt. Then all three CPUs read what the others have written, and compare it with their own memory status information. If they all agree, then the memory module is voted off-line. If not, and one CPU shows bad status for a memory module but the others show good status, then that CPU is voted off-line.

Fault-Tolerant Power Supply:

Referring now to Figure 22, the system of the preferred embodiment may use a fault-tolerant power supply which provides the capability for on-line replacement of failed power supply modules, as well as on-line replacement of CPU modules, memory modules, I/O processor modules, I/O controllers and disk modules as discussed above. In the circuit of Figure 22, an a/c power line 160 is connected directly to a power distribution unit 161 that provides power line filtering, transient suppressors, and a circuit breaker to protect against short circuits. To protect against a/c power line failure, redundant battery packs 162 and 163 provide 4-1/2 minutes of turn system power so that orderly system shutdown can be accomplished. Only one of the two battery packs 162 or 163 is required to be operative to safely shut the system down.

The power subsystem has two identical AC to DC bulk power supplies 164 and 165 which exhibit high power factor and energize a pair of 36-volt DC distribution busses 166 and 167. The system can

remain operational with one of the bulk power supplies 164 or 165 operational.

Four separate power distribution busses are included in these busses 166 and 167. The bulk supply 164 drives a power bus 166-1, 167-1, while the bulk supply 165 drives power bus 166-2, 167-2. The battery pack 162 drives bus 166-3, 167-3, and is itself recharged from both 166-1 and 166-2. The battery pack 163 drives bus 166-3, 167-3 and is recharged from busses 166-1 and 167-2. The three CPUs 11, 12 and 13 are driven from different combinations of these four distribution busses.

A number of DC-to-DC converters 168 connected to these 36-v busses 166 and 167 are used to individually power the CPU modules 11, 12 and 13, the memory modules 14 and 15, the I/O processors 26 and 27, and the I/O controllers 30. The bulk power supplies 164 and 165 also power the three system fans 169, and battery chargers for the battery packs 162 and 163. By having these separate DC-to-DC converters for each system component, failure of one converter does not result in system shutdown, but instead the system will continue under one of its failure recovery modes discussed above, and the failed power supply component can be replaced while the system is operating.

The power system can be shut down by either a manual switch (with standby and off functions) or under software control from a maintenance and diagnostic processor 170 which automatically defaults to the power-on state in the event of a maintenance and diagnostic power failure.

Fig. 23 is a block diagram of a data processing system according to another embodiment of the present invention. As shown therein, a data processing system comprises a CPU-A, a CPU-B, and a CPU-C which communicate with a memory A, a memory B, and a memory C through a CPU-memory bus AA, a CPU-memory bus BB, and a CPU-memory bus CC, respectively. CPU-A, CPU-B, and CPU-C communicate data to an output interface OI through buses AO, BO, and CO, respectively, and they receive data from an input interface II through buses IA, IB, and IC, respectively. CPU-A, CPU-B, and CPU-C communicate with each other and to output interface 28 and input interface 40 through a sync bus ABC. Output interface OI communicates data from CPU-A, CPU-E and CPU-C to a vote circuit V over an interface-vote bus VB. Vote circuit V determines the data from which processor should be communicated to an I/O controller IOC. Data is communicated to I/O controller IOC and ultimately to an I/O device IOD over a vote-controller bus VC and a controller-device bus CD, respectively. Data is communicated from I/O controller IOC to input interface II through an interface-controller bus ICB.

As shown in Fig. 24, CPU-A, CPU-E, and CPU-C execute a number of instructions, some of which have processor events associated with them. A processor event is defined either explicitly or implicitly by the code running in the processors. For example, each microprocessor write operation may be considered a processor event. Other possibilities for a processor event may be data reads, bus transfers, or specific signals generated by explicit code in the processor. (In the embodiment of Figs. 1-22, an "event" is a Run cycle, i.e., a machine or clock cycle of the CPU where the pipeline advances so a stall is not in effect.) In any case, events occur in the same order in each processor. However, events may not occur at the same time in each processor. For example, one processor may be executing a revised version of the code executing in another processor, and the revised code may contain extra instructions which cause the events to occur at different times. Notice event-4 in CPU-A and CPU-B. Another reason events may not occur at the same time, even in programs executing identical code, is the occurrence of unexpected errors such as cache misses which necessitate read retries or the detection of parity errors in which case execution may branch to an error routine. Notice the errors encountered by CPU-A and CPU-C.

Because a system according to this embodiment encounters the plurality of processor events in the same order, the processors may be synchronized by synchronizing each microprocessor to a prescribed event. The structure which allows each processor to be so synchronized is illustrated for CPU-A in Fig. 25. CPU-B and CPU-C are constructed the same way. As shown in Fig. 25, CPU-A comprises a processor 180 which executes instructions based on clock pulses received from its own clock 184 over a line 186. Processor 180 generates an internal sync request signal on a line 190, a clock output signal on a line 191, an "extra" clock signal on a line 192, and a processor event signal on a line 194. "Extra" clock cycles may occur because of error retries, variations in cache hit rates, asynchronous logic or for other reasons. They represent clock cycles which ordinarily do not occur when executing a particular program. Processor 180 receives a wait signal on a line 196 and an interrupt signal on a line 198.

CPU-A further includes a cycle counter 200 for counting clock cycles, an event counter 202 for counting processor events, a compare circuit 206 for comparing the value of event counter 202 with the other event counters within the system, and a sync logic circuit 210 for controlling the synchronization of CPU A. Cycle counter 200 is connected to line 191 for counting the number of clock cycles occurring since the last processor event. Cycle

counter 200 is connected to line 192 through an inverter 214 for inhibiting clock cycle counting when extra clock cycles occur. The reason for this is discussed below. Cycle counter 200 also is connected to line 194 for being reset upon each processor event. Whenever cycle counter 200 overflows, it generates an interrupt signal on a line 218 which, in turn, is connected to an OR gate 222. The output of OR gate 222 is connected to line 198. It is understood that OR gate 222 is a conceptual OR gate. Actual interrupt processing is performed using well known techniques.

Event counter 202 is connected to line 194 for counting events detected by processor 180. As with cycle counter 200, whenever event counter 202 overflows it generates an interrupt signal to OR gate 222 on a line 226. The value of event counter 202 is communicated to compare circuit 206 over a line 228 and to sync bus ABC over line 230. Compare circuit 206 receives the number of events counted by event counter 202 over line 230 and the number of events counted by the event counters for the other processors within the system from sync bus ABC. Compare circuit 206 generates a signal to sync logic circuit 210 over a line 234 indicating the relationship between the value from event counter 202 and the values from the other event counters in the system.

Sync logic circuit 210 controls the operation of processor 180 by asserting or removing a wait signal on line 196 in response to events indicated on line 194. Then the processors are synchronized, sync logic circuit 210 generates synchronized external interrupt signals on line 238.

Operation of the system may be understood by referring to Figs. 26A-26C and 27.

Fig. 26A illustrates a situation where a sync request (e.g., external interrupt) is received by CPUs-A, -B and -C, but where CPUs -A, -B and -C are executing different portions of code. In this case, CPU-A executes code until event-4 is indicated in time at a point 290. When event-4 is detected, sync logic circuit 210 causes CPU-A to enter a wait state. Similarly, CPU-B continues running until event-5 is detected at a point 251 whereupon CPU-B enters a wait state. CPU-C executes code until event-6 is detected at a point 252 whereupon CPU-C enters a wait state. Since the number of events counted by CPU-A is less than the number of events counted by CPU-C, CPU-A resumes instruction execution at a point 253 until event-5 is detected at a point 254 whereupon CPU -A again enters a wait state. When it is ascertained that CPU-A still is behind CPU-C, instruction execution resumes at a point 255 until event-6 is detected at a point 256 whereupon CPU-A again enters a wait state. CPU-B undergoes a similar processing sequence. That is, when it is ascertained that the

number of events counted for CPU-B is less than the maximum number of events counted by any of the three processors, CPU-B resumes instruction execution at a point 257 until the next event is detected at point 258 whereupon CPU-B enters a wait state, and so on. Since CPU-C counted the most events before the sync request was received, CPU-C remains in a wait state until the event counters for each of CPU-A, -B and -C are equal. When this occurs, the sync logic circuit 210 associated with each processor issues a synchronized external interrupt signal on line 238, releases the wait signal on line 196 and execution for each processor resumes at a common point 259.

Fig. 26B illustrates a processing sequence wherein CPUs-A, -B and -C are synchronized as a result of an event counter overflow. As shown therein, CPU-A executes code and an event occurs at a point 260. Code execution resumes until the event counter for CPU-A overflows at a point 261. At this point, event counter 202 in CPU-A issues an interrupt signal on line 226, and CPU-A enters a wait state. The same sequence of events and event counter overflows occur for CPUs-B and -C at points 262, 263 and at points 264, 265 respectively. When it is ascertained that each CPU-A, -B and -C is in a wait state, sync logic circuit 210 for each processor removes the signal from line 196, and code execution resumes at common points 266.

Fig. 26C illustrates a situation where synchronization occurs as a result of a cycle counter overflow. As shown therein, CPU-A detects event-7 at a point 267, and code execution continues for $2^{n(\text{cycle counter})}$ clock cycles until its cycle counter 200 overflows at a point 268. Cycle counter 200 issues an interrupt signal on line 218, and CPU-A enters a wait state. Similarly, CPU-E detects event-7 at a point 269 and continues code execution until its cycle counter overflows at a point 270 whereupon CPU-B enters a wait state. Finally CPU-C detects event-7 at a point 271 and continues code execution until its cycle counter 200 overflows at a point 272 whereupon CPU-C enters a wait state. When it is ascertained that each processor is in a wait state, sync logic circuit 210 removes the signal from line 196, and code execution resumes at common points 273.

Fig. 27 illustrates the processing sequence for each CPU-A, -B and -C. As shown therein, each processor is clocked for executing instructions in a step 300. If it is ascertained in a step 304 that the present clock cycle is an extra clock cycle, then cycle counter 200 is inhibited in a step 308, and processing resumes in step 300. If it is ascertained in step 304 that the present clock cycle is not an extra clock cycle, then cycle counter 200 is incremented in a step 312. It is then ascertained in a step 316 whether cycle counter 200 has over-

flowed. If so, then cycle counter 200 interrupts its respective processor in a step 320 and generates a sync request in step 324. In response to the interrupt generated by cycle counter 200, the processor generates an event in a step 328 whereupon the processor enters a wait state in a step 332 (as a result of the sync request signal generated in step 324).

In the case of a cycle counter overflow, all processors reach the processor event in the interrupt code exactly $2^{n(\text{cycle counter})}$ clock cycles after the last processor event. The processors will be in sync as long as not extra clock cycles occurred during the $2^{n(\text{cycle counter})}$ clock cycles before the cycle counter overflows. If extra clocks did occur, the processors would stop at different points, possibly causing a mismatch, and the processors may respond differently to the interrupts, thus resulting in a processor or system failure. That is why each cycle counter is disabled during every extra clock cycle.

If it is ascertained in 316 that cycle counter 200 did not overflow, it is then ascertained in a step 340 whether an event has occurred. If not, processing resumes in step 300. If an event has occurred, then event counter 202 is incremented in a step 344. It is then ascertained in a step 348 whether event counter 202 has overflowed. If so, then the processor is interrupted in step 320 and processing continues as with a cycle counter overflow until the processor is halted in step 332. If it is ascertained in step 348 that event counter 202 has not overflowed, then it is ascertained in a step 352 whether a sync request is outstanding. If so, then processor is halted in step 332; otherwise code execution continues in step 300.

After the processor is stopped in step 332, it is then ascertained in a step 360 whether all processors have stopped. If not, then it is ascertained in a step 364 whether the maximum amount of time allowed for processor synchronization has been exceeded. This may occur, for example, upon a processor failure. If the maximum time has been exceeded, then the processor is voted out, i.e., disregarded in the comparison process, in a step 368. In any event, comparison continues among the properly functioning processors in step 360. Once all processors have stopped, the counters are compared in a step 372. If it is ascertained that the counter for a particular processor is greater than another processor, then the processor(s) having the greater count remain(s) in a wait state and processing continues in step 360. On the other hand, if the value of the counter for a particular processor is less than the maximum count value, the processing reverts to step 300 wherein execution resumes until the next event is detected, the processor stopped, and the counters are again compared. If it is ascer-

tained in step 372 that all counters are equal, then the wait signals are removed from each processor, and the processors are restarted in a step 376 for servicing the sync request.

While the above is a complete description the embodiment of Fig. 23-27 of the present invention, various modifications may be employed. For example, sync logic circuit 210 may be integrated into a single circuit connected to each chip, and the system may be used with any number of processors.

The embodiment of Figs. 23-27 shows a method and apparatus for synchronizing a plurality of processors. Each processor runs off its own independent clock, indicates the occurrence of a prescribed process or event on one line and receives signals on another line for initiating a processor wait state. Each processor has a counter which counts the number of processor events indicated since the last time the processors were synchronized. When an event requiring synchronization is detected by a sync logic circuit associated with the processor, the sync logic circuit generates the wait signal after the next processor event. A compare circuit associated with each processor then tests the other event counters in the system and determines whether its associated processor is behind the others. If so, the sync logic circuit removes the wait signal until the next processor event. The processor is finally stopped when its event counter matches the event counter for the fastest processor. At that time, all processors are synchronized and may be restarted for servicing the event. If no synchronizing event occurs before an event counter reaches its maximum value, and overflow of the event counter forces resynchronization, a cycle counter is provided for counting the number of clock cycles since the last processor event. The cycle counter is set to overflow and force resynchronization at a point before maximum interrupt latency time is exceeded.

An additional embodiment of the invention is illustrated in Figure 28, where a fault-tolerant computer system is shown having two or more identical CPU modules CPU-A and CPU-B, each with its own memory 16, and each having its own clock oscillator 17. These processors are loosely synchronized via bus 18 according to the features of my prior application Ser. No. 118,503, using counter circuits, although any synchronization mechanism to keep the processor drift below some limit could be used.

In Figure 28, each of the processors CPU-A and CPU-B is connected by a bus 21 or 22 to data output and input modules 14 and 15; output data from the processors is voted by vote circuits 100 in these modules 14 and 15. A FIFO type of output buffer 50 may be employed in each processor to queue up outbound data. Since, in this illustration,

the memory 16 is in the CPU module, the outbound data will be I/O output data. Thus, outbound data is connected from the output of buffer 50 in CPU-A to the voter 100 of module 14 by a bus 21a and to voter 100 of module 15 by a bus 21b, and likewise outbound data is connected from CPU-B to voters in modules 14 and 15 by busses 22a and 22b. Incoming data to each CPU module is on busses 21c and 21d in bus 21, or on busses 22c and 22d in bus 22. Note that the incoming data is not voted.

The voters 100 of Figure 28, in the case of three or more CPU modules, takes a majority vote of the output data from the CPUs. In the case of two processors, the voters detect the differences between the two outputs, and if they are identical the voter passes one copy of the identical data through to I/O busses 24 and 25. The voter waits until all CPUs have issued an data output on their busses 21 or 22 before a vote is performed. The processors, if executing the same code, and within the limiting bounds of the loose synchronization mechanism, will send the same data in the same order (from all non-faulty processors). Note that only one set of voters 100 is required. If a processor malfunctions and begins to write incorrect data into its memory 16, that malfunction is of interest and is detected only if the bad data is sent to the outside world (displays, keyboard, printer, co-processors, etc., on the I/O busses), or if the malfunction causes the processor to get out of synch with the other processors.

Incoming data from the I/O busses is connected from the busses 24 and 25 through buffers 21e and 21f to the input busses 21c and 21d, and likewise through buffers 22e and 22f to incoming busses 22c and 22d, in the embodiment of Figure 28. On inbound I/O, data is loaded to all these buffers, then the CPUs asynchronously unload these buffers. These buffers may buffer a single byte, word or packet, or they may be FIFOs or circular buffers which hold multiple data items and allow simultaneous loading and emptying. In the simplest case each buffer 21e, 21f, 22e and 22f is only a single data register with an empty/full flag, so new data would not be loaded from the I/O busses 24 and 25 until all processors had signalled empty (or a timeout occurred). Multiple I/O busses could share a buffer, or have independent buffers.

To be more fault-secure, the I/O busses of Figure 28 preferably have parity checking or code checking to allow each CPU to detect incorrect I/O data before loading it to memory. Otherwise, a bad I/O bus could cause the memories of all processors to be corrupted.

The embodiment of Figure 28 thus shows a fault-tolerant computer system employs multiple identical CPUs executing the same instruction

stream, each with their own independent memory. The multiple CPUs are loosely synchronized, as by counting events such as operating cycles and stalling any CPU ahead of others. Data output references via separate busses are voted at separate ports of each of the CPUs by voting circuits which detect when all CPUs have made the same reference, and only then pass on identical references to external I/O busses. The ports may include FIFO buffers to allow output references from the asynchronous CPUs to be handled as the CPUs load the FIFOs at different times. Input data to the CPUs from the I/O busses is not voted, but is buffered to allow the CPUs to accept it at their own clock rate.

While the invention has been described with reference to specific embodiments, the description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiment, as well as other embodiments of the invention, will be apparent to persons skilled in the art upon reference to this description. It is therefore contemplated that the appended claims will cover any such modifications or embodiments as fall within the true scope of the invention.

Claims

1. A multiple CPU system, comprising:
 - a) a plurality of CPUs executing an instruction stream, the CPUs each being clocked independently of one another to provide separate machine cycles for each CPU, said machine cycles including execution cycles where an instruction of said instruction stream is executed and stall cycles where an instruction of said instruction stream is not executed, each CPU having a memory request input/output port;
 - b) a common memory coupled to the input/output ports of said CPUs, the common memory implementing a memory request only after receiving the same request from all of said CPUs, the memory sending an acknowledge signal to the CPUs when implementing a memory request, each of the CPUs executing stall cycles while awaiting implementation of a memory request by the common memory as signalled by said acknowledge signal;
 - c) each of the CPUs having a counter to count execution cycles but not stall cycles; and
 - d) said CPUs having an interrupt circuit responsive to an external interrupt request and coupled to said counters in said CPUs and responsive to a selected count in each of said counters for separately interrupting each CPU at the same execution cycle
2. A system according to claim 1 wherein each of said CPUs has a local memory not accessible by the other CPUs.
3. A multiple CPU system with synchronization of external interrupts, comprising:
 - a) a plurality of CPUs independently executing the same instruction stream, the CPUs each being clocked independently of one another to provide execution cycles during which instructions of said instruction stream are executed and to provide stall cycles during which instructions are not executed;
 - b) each of the CPUs having a counter to count said execution cycles but not stall cycles;
 - c) and an interrupt circuit connected to all of said CPUs and responsive to an external interrupt request, said interrupt circuit being responsive to a selected count in each of said counters for interrupting each CPU separately at the same execution cycle in said instruction stream.
4. A system according to claim 3 wherein there are three said CPUs, and wherein all three CPUs access a common memory module, requests to said common memory from said CPUs being voted by said memory module.
5. A multiple CPU system, comprising:
 - a) a plurality of CPUs each executing an instruction stream, the CPUs being clocked independently of one another to provide execution cycles,
 - b) each of the CPUs having a counter to count execution cycles;
 - c) and an interrupt circuit connected to each of said CPUs and responsive to an external interrupt request, said interrupt circuit being responsive to a selected count in each of said counters for separately interrupting each CPU at the same execution cycle in said instruction stream.
6. A system according to claim 5 wherein there are three said CPUs, and wherein all three CPUs access a common memory module, requests to said common memory from said CPUs being voted by said memory module.
7. A multiple CPU system, comprising:
 - a) a plurality of CPUs, each of the CPUs independently executing the same instruction stream, the CPUs being clocked in-

dependently of one another to provide execution cycles, the CPUs each having an input/output port, at least one shared input/output device being coupled to said input/output ports of the plurality of CPUs;

b) each of the CPUs having a modulo N counter to count execution cycles;

c) and an interrupt circuit coupled to each of said CPUs and responsive to an external interrupt request, said interrupt circuit being correlated with said counters for applying an interrupt separately to each CPU at a preset count value of each of the counters, whereby each CPU is interrupted at the same execution cycle in said instruction stream.

8. A system according to claim 7 wherein there are three said CPUs, and wherein all three CPUs access a common memory module, requests to said common memory from said CPUs being voted by said memory module.

9. A computer system comprising:

a) a plurality of CPUs each executing an instruction stream, the CPUs being clocked independently of one another to define execution cycles;

b) each of the CPUs having counting means to count events related to execution cycles;

c) each CPU having a synchronizing circuit responsive to an externally-applied request applied to all of said CPUs, said synchronizing circuit for each CPU receiving input from all of said CPUs, for separately signalling each one of the CPUs at the same point in said instruction stream in response to said externally-applied request, the synchronizing circuit for each of said CPUs also responsive to said counting means indicating a selected maximum count and responsive to information from the other CPUs for causing the CPU to begin execution at the same execution cycle in said instruction stream.

10. A system according to claim 9 wherein said synchronizing request is an external interrupt.

11. A system according to claim 1 wherein there are three of said CPUs, and wherein said requests to said common memory via said port are voted by said common memory.

12. A system according to claim 1 wherein said interrupt circuit interrupts said CPUs only on a selected value registered by said counters.

13. A system according to claim 1 wherein said

interrupt circuit includes for each CPU means responsive to an indication of receipt of an interrupt request when a selected value is registered in said counter of each of the other CPUs.

14. A system according to claim 3 wherein each of said CPUs has a local memory not accessible by the other CPUs.

15. A system according to claim 4 wherein there are three of said CPUs, and a common memory accessed by said three CPUs, and accesses to said common memory by the CPUs are voted by said common memory.

16. A system according to claim 5 wherein each of said CPUs has a local memory not accessible by the other CPUs.

17. A system according to claim 5 wherein there are three of said CPUs, and wherein each of said CPUs makes access requests to a common memory, and said access requests are voted by said common memory.

18. A system according to claim 5 wherein said interrupt circuit includes for each CPU means responsive to receipt of an interrupt request when a selected value is registered in said counter of each of the other CPUs.

19. A system according to claim 7 wherein said shared input/output device is a common memory.

20. A system according to claim 7 wherein said interrupt circuit interrupts each of said CPUs when a selected value is registered by said counter for each CPU without stalling execution of instructions by faster ones of said CPUs.

21. A system according to claim 7 wherein said interrupt circuit includes for each CPU means responsive to receipt of an interrupt request when a selected value is registered in said counter of each of the other CPUs.

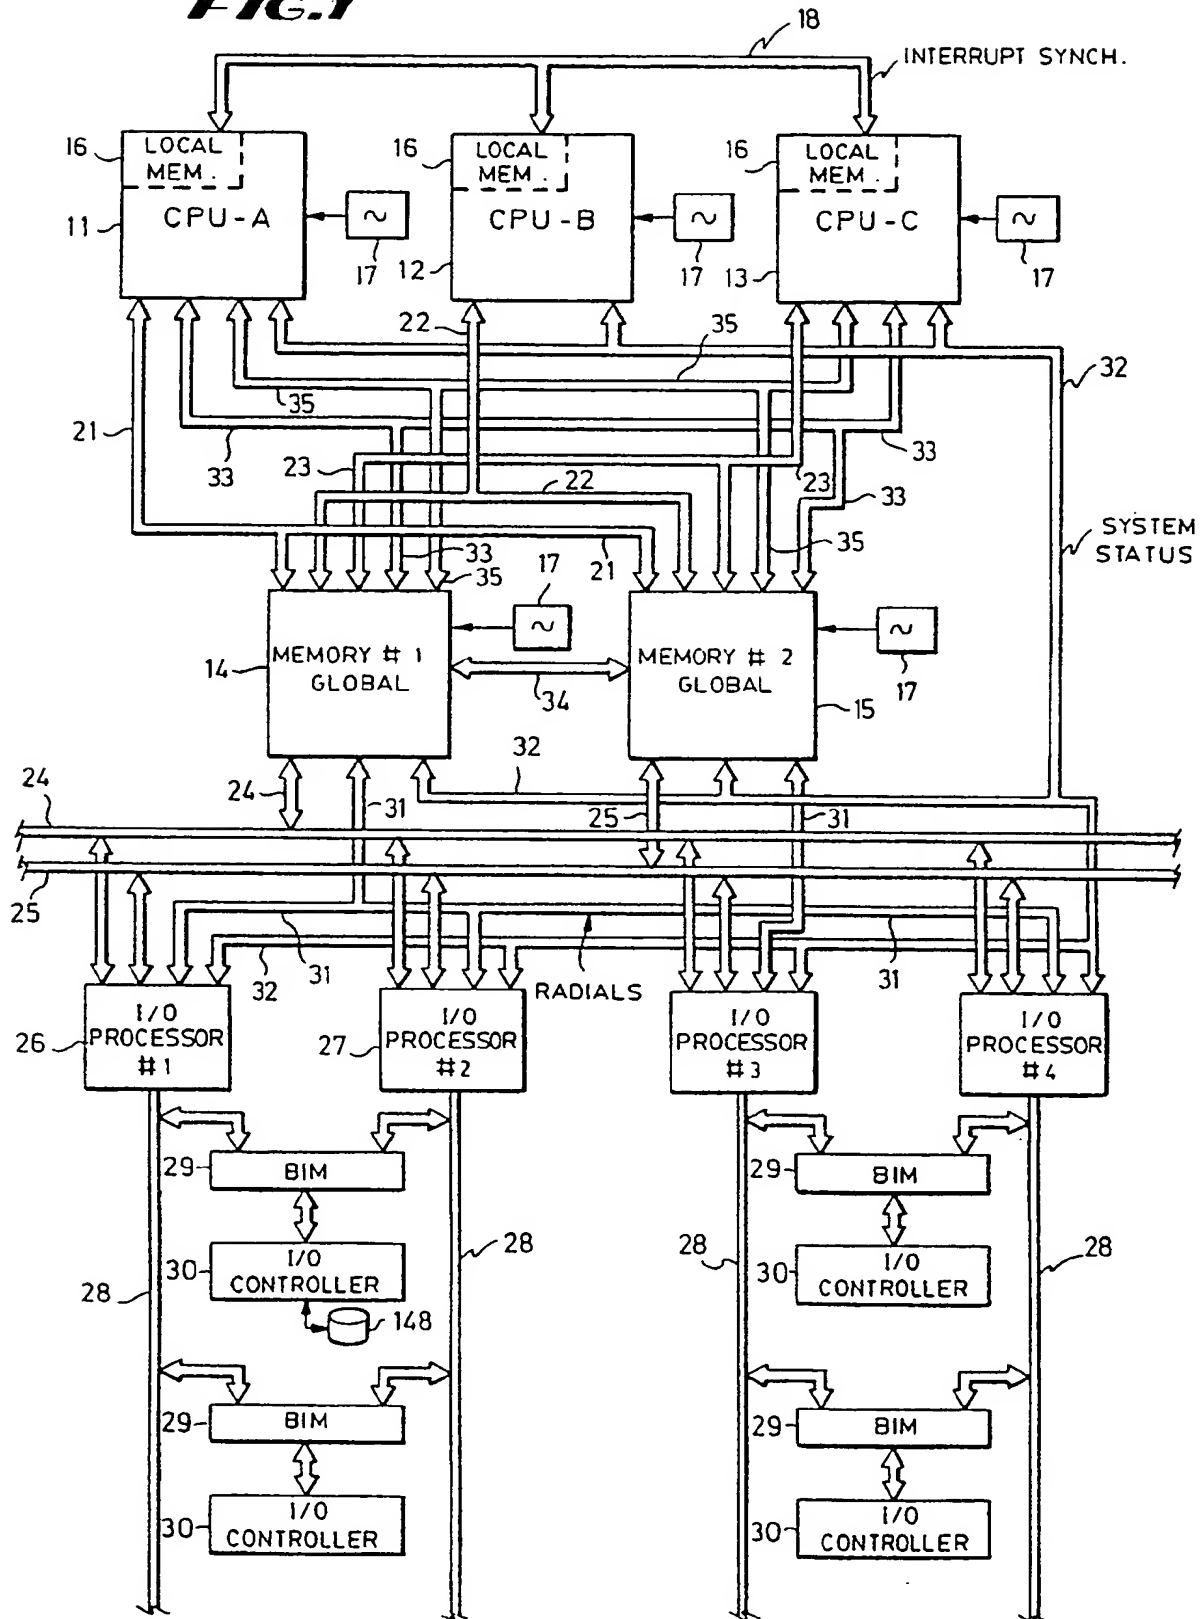
22. A system according to claim 7 wherein each of said CPUs has a local memory not accessible by the other CPUs.


23. A method of operating a computer system having a plurality of separately-clocked CPUs, comprising the steps of:

a) executing the same instruction stream on each of said CPUs;

- b) counting the instructions executed on each of said CPUs;
- c) detecting an external interrupt request and applying an interrupt signal separately to each one of the CPUs at a selected instruction count while the other of the CPUs continue to execute instructions.
24. A method according to claim 23 including the step of making memory access requests by said CPUs to a common memory and voting each of said requests by said common memory.
25. A method according to claim 24 including the step of accessing a local memory separately by each of said CPUs, the local memory for each CPU not being accessible by the other CPUs.
26. A method according to claim 23 wherein said step of detecting an external interrupt request is performed only at a selected value of said count of instructions.
27. Apparatus for synchronizing a plurality of processors, comprising:
 event detecting means in each one of said processors producing an indication of the occurrence of a selected type of event within the processor;
 event counting means responsive to said indication for counting the number of events for each one of said processors; and
 means responsive to said event counting means for altering processing of each one of said processors if the number of events counted for one of the processors is greater than the number of events counted for other of said processors.
28. Apparatus according to claim 27 wherein said selected type of circuit is a machine cycle in the processor in which the pipeline advances.
29. Apparatus according to claim 28 wherein said event counter is a cycle counter.
30. Apparatus according to claim 27 including sync-request means receiving a synchronization request signal, and wherein the means for suspending suspends processing of a processor in response to the synchronization request signal when the number of events counted for that processor is not less than the number of events counted for another processor.
31. Apparatus according to claim 30 wherein said synchronization request signal is an interrupt.
32. Apparatus according to claim 28 including means for restarting a suspended processor when the number of events counted for each processor is equal.
33. Apparatus according to claim 29 wherein the means for suspending suspends processing of a processor when the number of events counted for that processor is not less than the number of events counted for a suspended processor.
34. Apparatus for synchronizing a plurality of processors comprising, for each processor:
 event counting means, connected to the processor for counting the number of occurrences of a prescribed event;
 comparison means connected for receiving signals for the event counting means for each processor; and
 synchronization means, connected to receive a sync request input signal and to the event counter, and responsive to said comparison means, for suspending processing of a processor in response to a synchronization request signal until the number of events counted to each processor is equal to the number of events counted for other processors.
35. Apparatus according to claim 34 wherein the prescribed event is a clock cycle of the processor in which an instruction is executed, and wherein said clock cycle is one in which the pipeline of the processor advances.
36. Apparatus according to claim 34 wherein the sync request input is an interrupt, and wherein the event counting means is a cycle counter.
37. A fault-tolerant computer system comprising:
 a) multiple processors, each executing the same instruction stream, each processor having an independent clock, and each processor having a memory independent of the other processors;
 b) a plurality of vote circuits, each one of the vote circuits separately receiving output data from each one of the processors, and producing a voter output to I/O means only when multiple processors have sent the same data output.
38. A system according to claim 37 wherein said processors are loosely synchronized by counting cycles of operation of the processors and stalling processors ahead of others.

39. A system according to claim 37 including input buffers coupling each said I/O means to each one of said processors, the processors accepting data from said input buffers asynchronously.
40. A system according to claim 37 wherein said output data from each processor is held in an output buffer until all said processor have loaded output data to their output buffer, and wherein said output buffer is a FIFO.
41. A method of operating a multiple processor system comprising the steps of:
- a) clocking each of said processors independent of one another;
 - b) executing the same instruction stream in each one of the processors;
 - c) storing data for each processor in a separate memory not accessible by the other processors;
 - d) presenting output data to an output port of each processor;
 - e) detecting the output data in all said ports in a vote circuit and voting said output data to pass on the output data that is the same from multiple processors to I/O means.
42. A method according to claim 41 including the step of storing said output data in a buffer at each one of said ports, and wherein each one of said buffers is a FIFO.
43. A method according to claim 42 including the step of loosely synchronizing said processors by counting cycles of operation of the processors and stalling processors ahead of others.
44. A method of operating a computer system comprising the steps of:
- a) executing the same instruction stream in at least first and second processors;
 - b) generating remote accesses in each of said first and second processors, the remote accesses being directed to separate first and second access ports;
 - c) detecting each one of said remote accesses at said first and second access ports, waiting until a remote access is detected at both the first and second access ports, then voting said remote accesses and passing along said remote accesses if both are the same.
45. A method according to claim 44 wherein said step of executing is in independently-clocked processors, and wherein said step of generating remote accesses in said first and second processors includes temporarily storing said remote accesses in a buffer for each processor.
46. A method according to claim 44 wherein said first and second access ports are in first and second modules operated asynchronous to said first and second processors, and wherein said first and second processors are loosely synchronized by counting cycles of operation and stalling a processor ahead of the other; and including the step of storing data for said first and second processors in separate memories accessible by only one processor.
47. A computer system comprising:
- a) first and second processors executing the same instruction stream;
 - b) means generating remote accesses in each of said first and second processors, the remote accesses being directed to separate first and second access ports;
 - c) separate voter means detecting each one of said remote accesses at said first and second access ports, the voter means waiting until a remote access is detected at both the first and second access ports before voting said remote accesses and passing along said remote accesses if both are the same.
48. A system according to claim 47 wherein said processors are independently clocked; and including a buffer for each processor temporarily storing said remote accesses.
49. A system according to claim 47 wherein said first and second access ports are in first and second modules operated asynchronous to said first and second processors.
50. A system according to claim 47 including separate memory means for each one of said first and second processors, each memory means accessible by only one processor.

FIG. 1



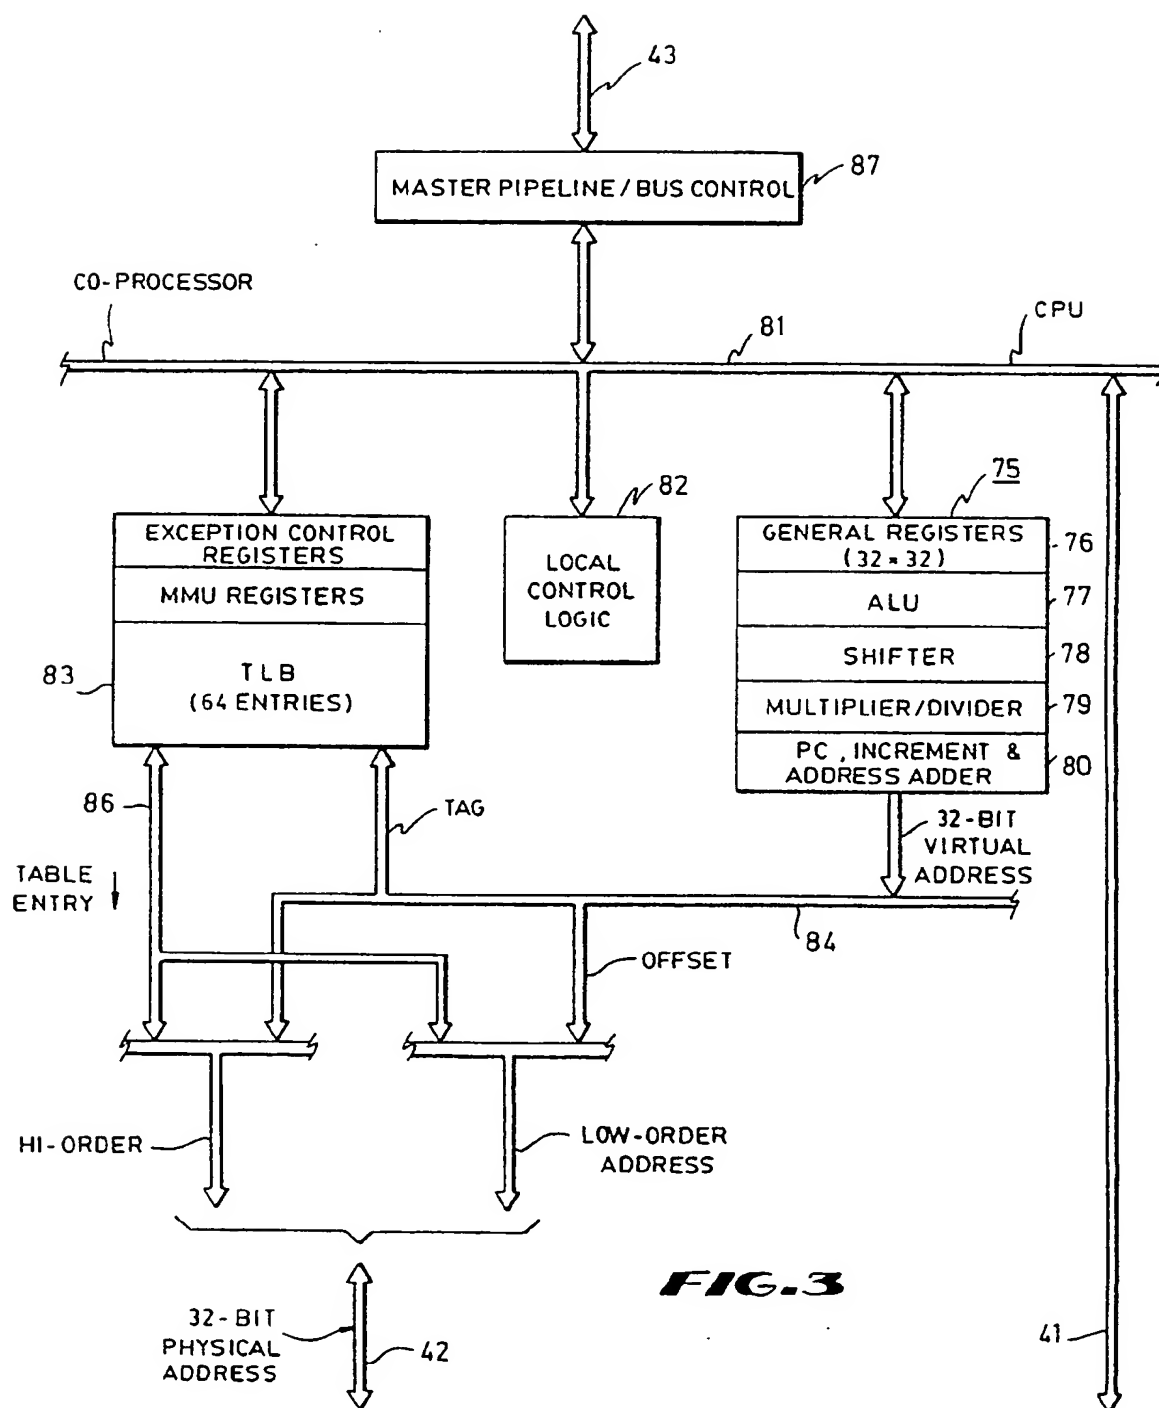


FIG. 3




FIG.4

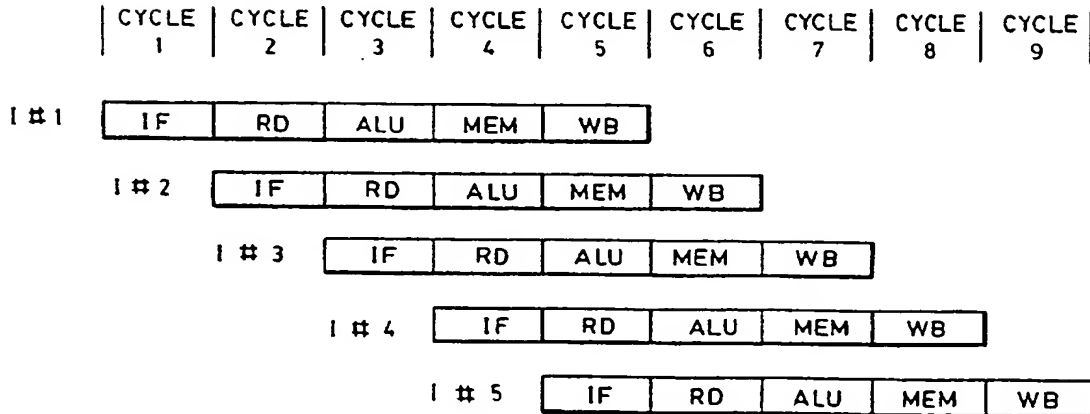


FIG.5


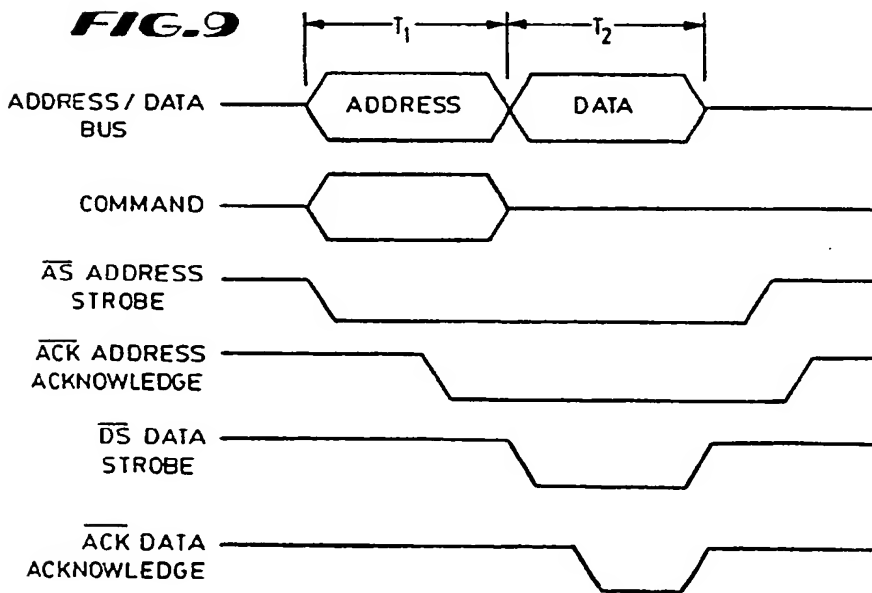



FIG. 6






FIG. 8









FIG.10a



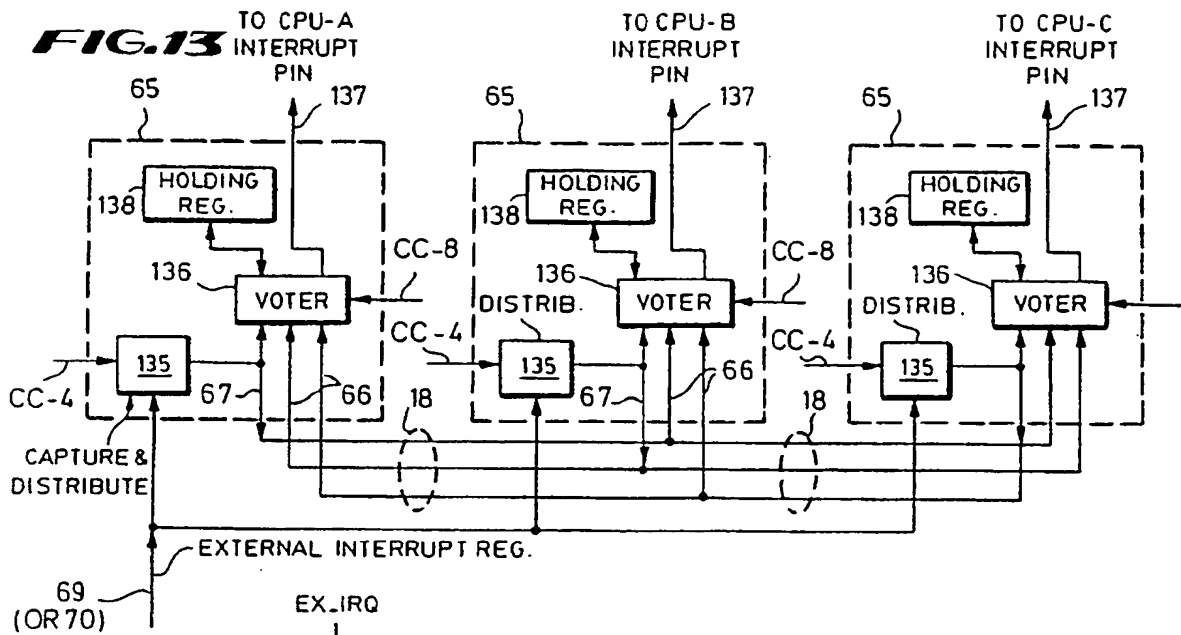


FIG.14

EX.IRQ

CPU A	0	1	2	3	4	5	6	7	8	9	10	11	12
CPU B	2	3	4	5	6	7	8	9	10	11	12	13	14
CPU C	3	4	5	6	7	8	9	10	11	12	13	14	15
IRQ.A.PENDING	0	0	0	0	1	1	1	1	1	1	1	1	1
IRQ.B.PENDING	0	0	1	1	1	1	1	1	1	1	1	1	1
IRQ.C.PENDING	0	1	1	1	1	1	1	1	1	1	1	1	1
IRQ.A	0	0	0	0	0	0	0	0	1	1	1	1	1
IRQ.B	0	0	0	0	0	0	1	1	1	1	1	1	1
IRQ.C	0	0	0	0	0	1	1	1	1	1	1	1	1
	t ₀	t ₁	t ₂	t ₃	t ₄	t ₅	t ₆	t ₇	t ₈	t ₉	t ₁₀	t ₁₁	t ₁₂

EX.IRQ

FIG.15

EX.IRQ

CPU A	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
CPU B	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1
CPU C	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2
EX.IRQ	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
IRQ.A.PENDING	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
IRQ.B.PENDING	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
IRQ.C.PENDING	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1
IRQ.A	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
IRQ.B	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1
IRQ.C	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1
	t ₀	t ₁	t ₂	t ₃	t ₄	t ₅	t ₆	t ₇	t ₈	t ₉	t ₁₀	t ₁₁	t ₁₂	t ₁₃	t ₁₄	t ₁₅

FIG. 18

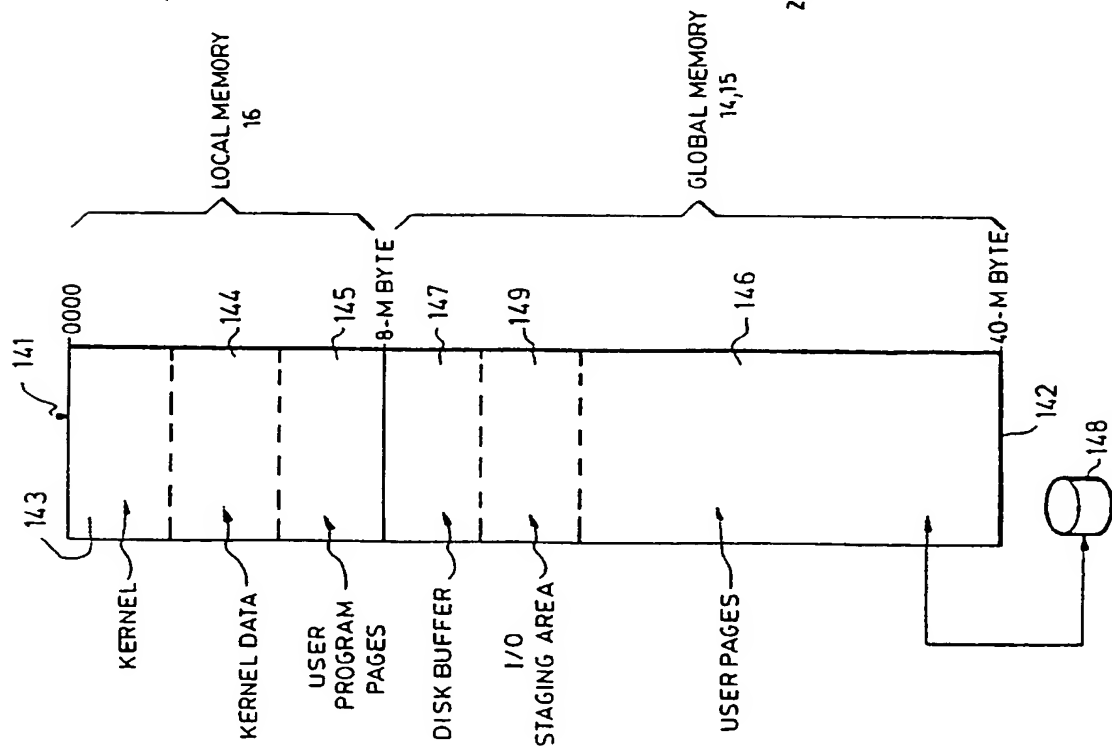



FIG. 19






FIG.20

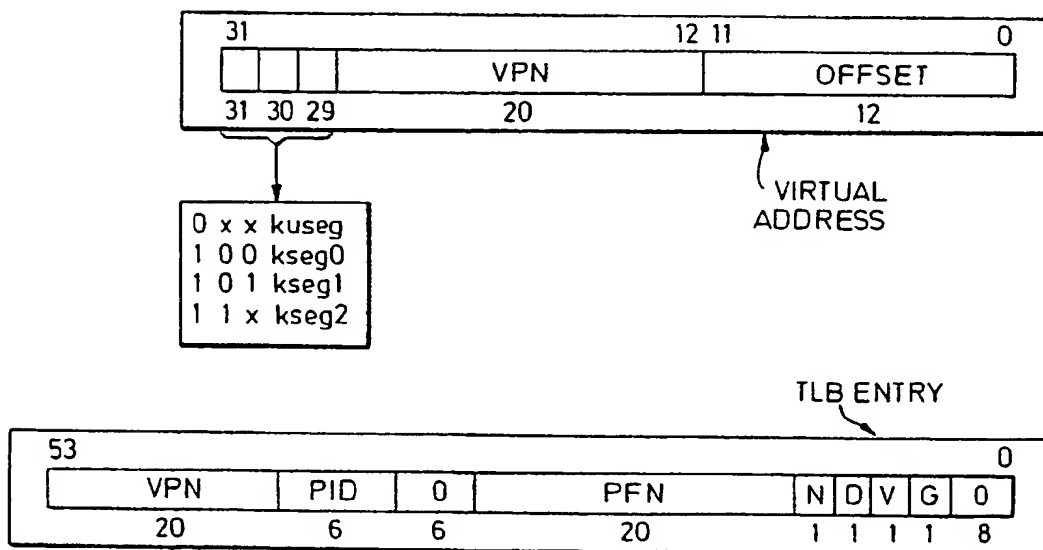
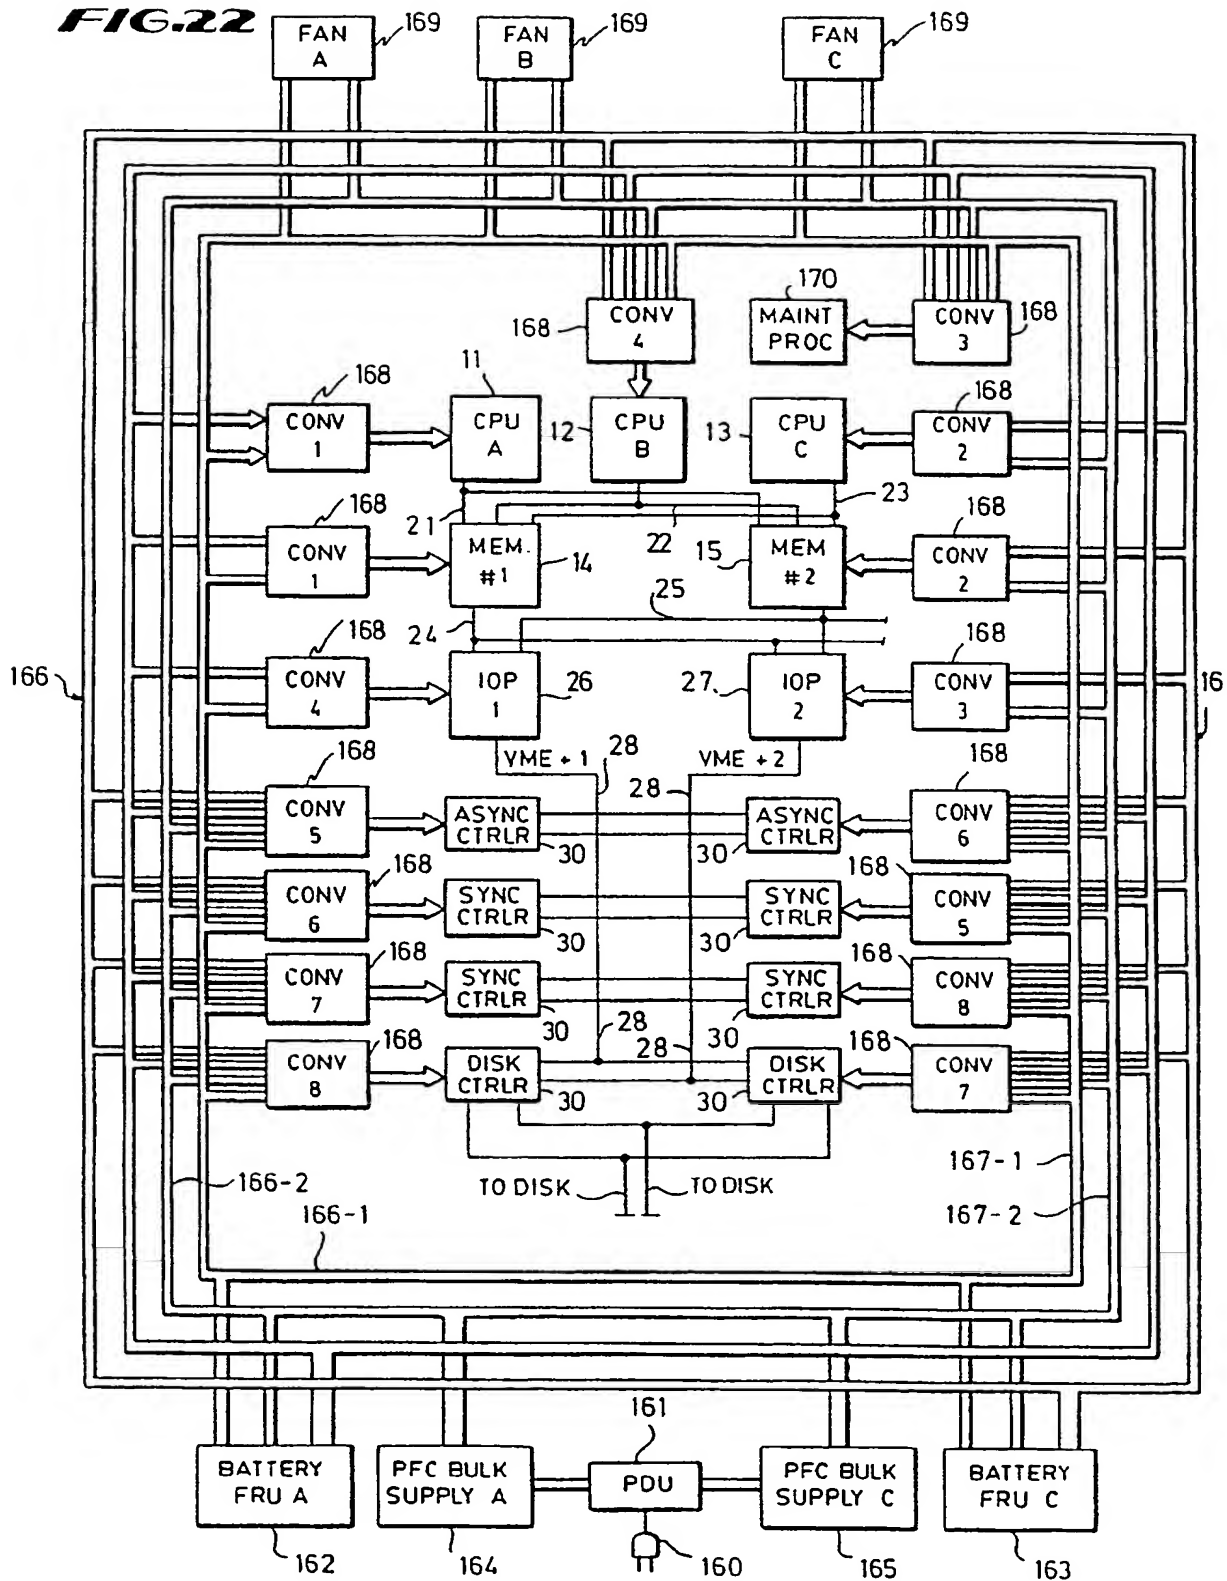


FIG. 22




FIG. 23

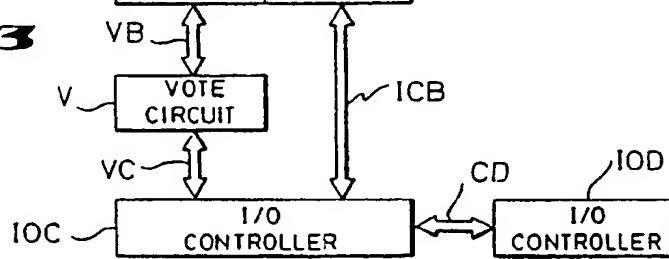


FIG. 24

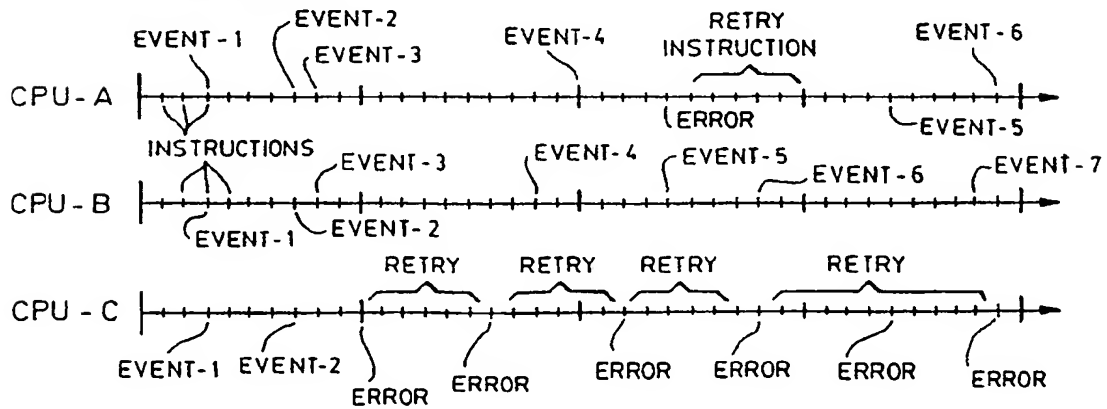



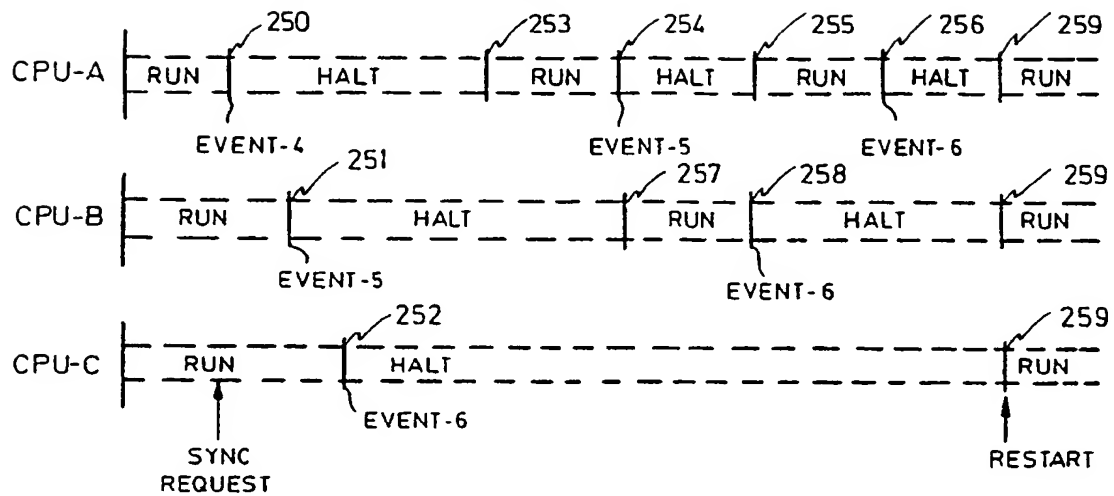
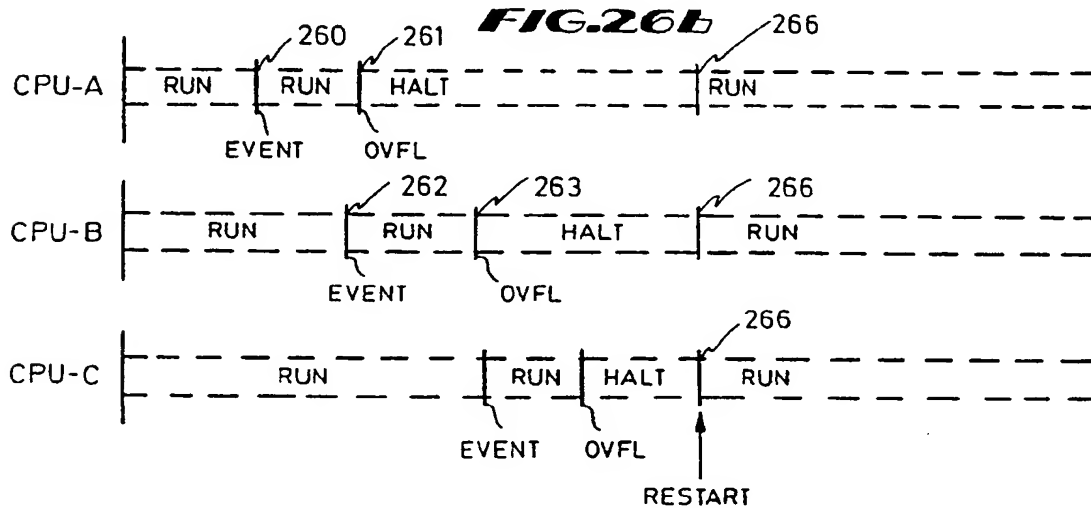
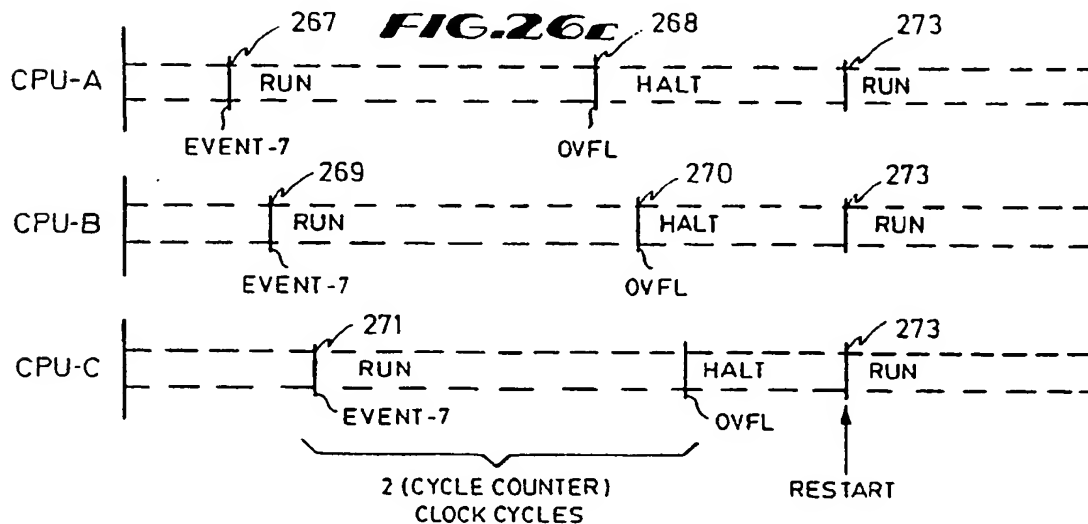
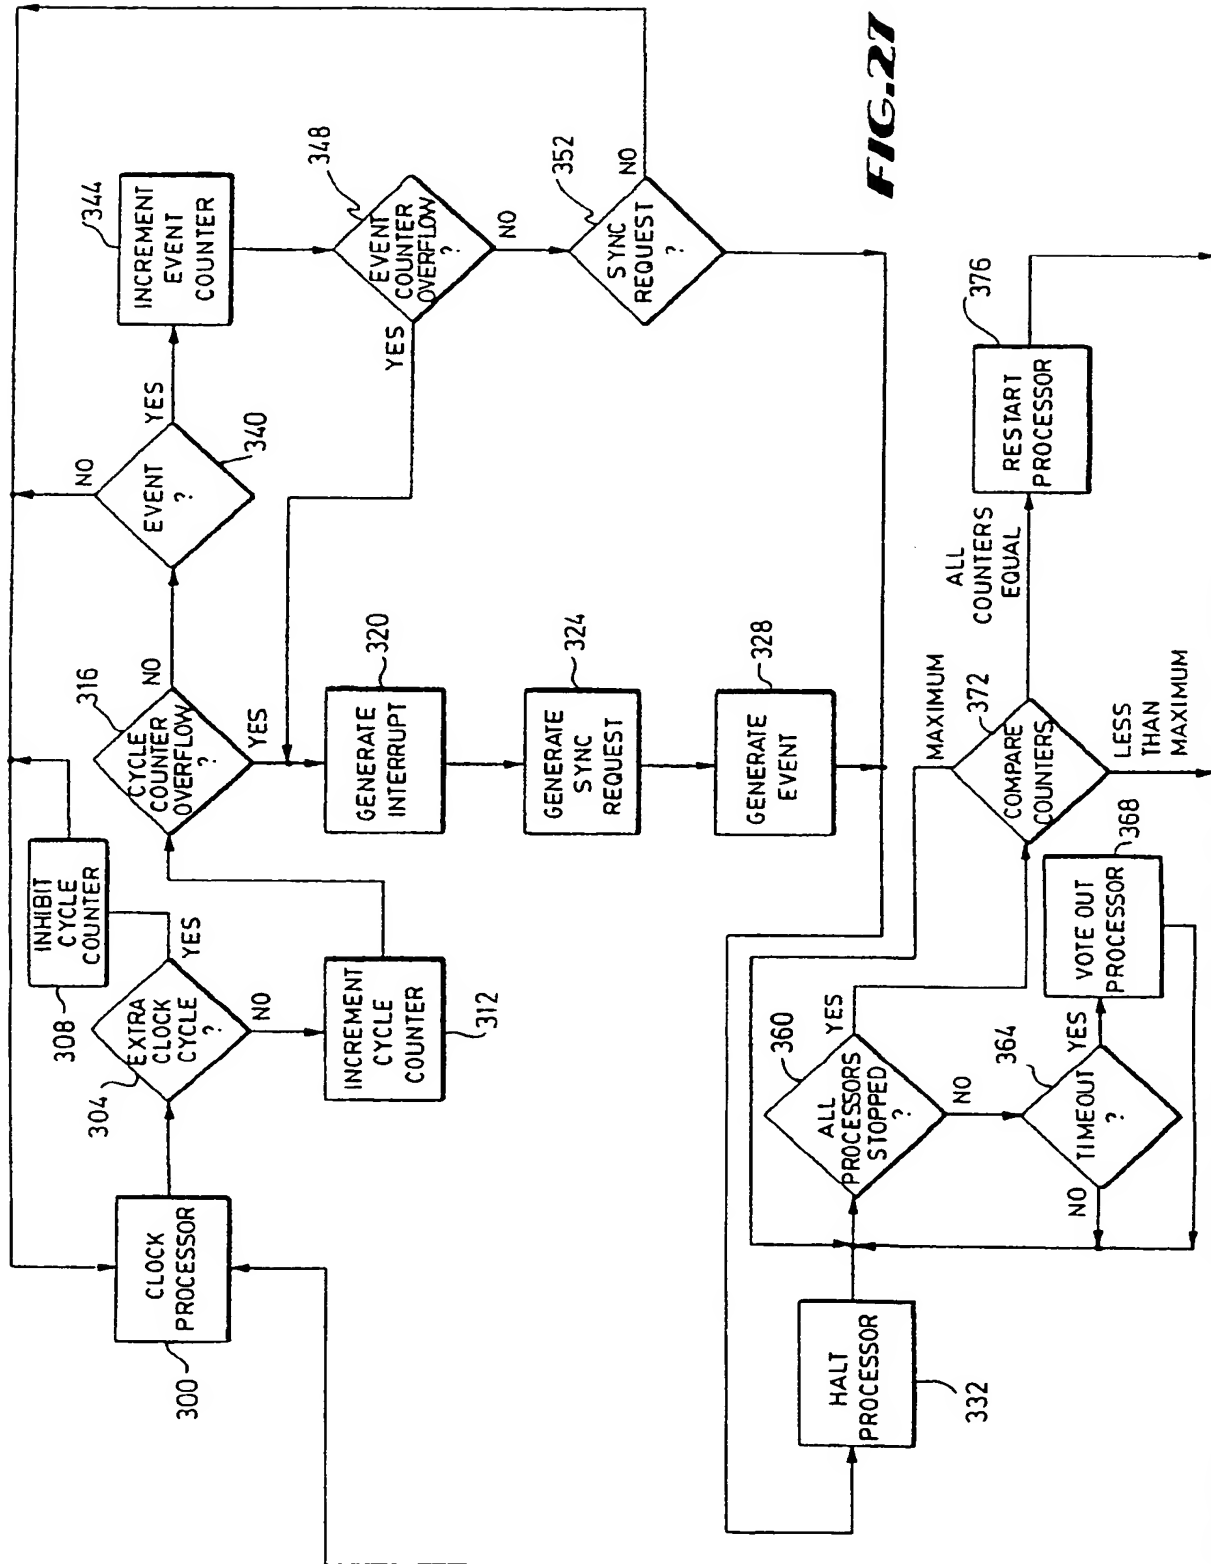
FIG.25

FIG.26a**FIG.26b****FIG.26c**






FIG.28



European
Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 90 10 5101

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
X	US-A-4 733 353 (V.C. JASWA) * Abstract; figure 1; column 3, lines 40-54 *	37,41	G 06 F 11/18
A		1,3-9,11, 15,17,19, 23,24, 27-29,32, 33,35 38-40, 42-49	
A			
A	COMPUTER DESIGN, vol. 21, no. 11, November 1982, pages 211-220; J.H. WENSLEY: "Fault tolerant systems can prevent timing problems" * Page 216, column 1, lines 12-21 *	1,3-9,11, 15,17,19, 23,24, 27-29,32, 33,35	
A	IDEM	37-49	
A	US-A-4 703 452 (R.J. ABRANT et al.) * Abstract; column 1, lines 10-24 *	3,5,7,9, 10,13,20, 24,26,30, 31,34,36	TECHNICAL FIELDS SEARCHED (Int. Cl.5)
A	EP-A-0 141 246 (SIEMENS AG) * Page 13, line 25 - page 14, line 9 *	2,14,16, 22,25,50	G 06 F
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of search 15 November 90	Examiner CORREMANS G.J.W.
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document			